CARNEGIE MELLON UNIVERSITY

School of Architecture
College of Fine Arts

Dissertation

Submitted in Partial Fulfillment of the requirements for the degree of

Doctor of Philosophy in Computational Design

TITLE:

Robotic Technologies and the Reconfiguration of Building Construction Work

AUTHOR:

Emek Erdolu

ACCEPTED BY ADVISORY COMMITTEE:

Daniel Cardoso Llach (Advisor) — September 23, 2024

Associate Professor, School of Architecture
Carnegie Mellon University — DATE

Daragh Byrne (Committee Member) — **September 27th, 2024**

Associate Teaching Professor, School of Architecture
Carnegie Mellon University — DATE

Molly Wright Steenson (External Committee Member) — September 26, 2024

President and CEO, American Swedish Institute
Associate Professor, School of Design
Carnegie Mellon University — DATE

Robotic Technologies and the Reconfiguration of Building Construction Work

Emek Erdolu
PhD Candidate in Computational Design, School of Architecture
Carnegie Mellon University

© 2024 Emek Erdolu. All Rights Reserved.

Carnegie Mellon University

September 2024

Dissertation Committee

Dr. Daniel Cardoso Llach (Advisor)
Associate Professor, School of Architecture
Computational Design Track Chair, School of Architecture
Carnegie Mellon University

Dr. Daragh Byrne (Committee Member)
Associate Teaching Professor, School of Architecture
Courtesy Appointment, School of Design and Human-Computer Interaction Institute
Carnegie Mellon University

Dr. Molly Wright Steenson (External Committee Member)
President and CEO, American Swedish Institute
Associate Professor, School of Design
Courtesy Appointment, School of Architecture
Carnegie Mellon University

Abstract

Advancements in robotics have led to the exploration of industrial arms, mobile robots, drones, and other semi-autonomous technologies in building construction. Drawing from primary historical and ethnographic research, this dissertation contributes a historically and socially grounded understanding of the field of construction robotics. It especially examines how technology designers conceptualize construction work in robotics practices, identifying a shift from high-level descriptions and classifications of construction tasks in the 1980s to situated engagements with workers and construction knowledge in the present. Revealing contrasts between early visions of automation and productivity with emergent approaches seeking adaptive robot support and worker–robot work-sharing, it traces the emergence of alternative systems, new skills, and kinds of interdisciplinary technical labor. Drawing from these findings, the dissertation develops three arguments. First, that technology designers' conceptions of work decisively shape system development in construction robotics. Second, that the participation of designers and workers in each other's knowledge worlds helps to conceive alternative systems and practically integrate robotic technologies into existing workflows. Lastly, that automation progresses faster in designers' discourses than in machines. From here, the dissertation outlines new considerations and pedagogies across computational design, human-computer interaction, robotics, and architecture for the sensible integration of these technologies in construction.

Acknowledgments

My deepest gratitude extends to my advisor Daniel Cardoso Llach for his care, guidance, and support. I thank him for showing me new intellectual horizons and encouraging me to sail towards. I would like to thank my committee member Daragh Byrne for his generosity and diligent involvement in many aspects of my research. I am also thankful to my committee member Molly Wright Steenson for her sharp comments and big smiles she offered each time.

Computational Design program provided me with the resources necessary to conduct this research. I am grateful to Ramesh Krishnamurti, who carried this program on his shoulders for many years, and Daniel Cardoso Llach and the current faculty for continuing to carry it. Aside from academic facilitation, the program brought friendships and many stimulating conversations to me over the years. I would like to assure each member that I will remember you.

In addition to my committee, I acknowledge a list of people for helping me acquire the knowledge key to this research: Katherine Barbera at Carnegie Mellon University Libraries, Henny Admoni at the Robotics Institute, and Gabriella Lukacs at the University of Pittsburgh Department of Anthropology. The list continues with Dominik Moritz, Adam Perer, Brad Myers, and Ken Koedinger at Carnegie Mellon's Human-Computer Interaction Institute. Thank you all.

Parts of this research would not be possible without my colleagues and participants. I am grateful to my long-term colleagues in the ReAC project, Yuning Wu and Jiaying Wei, for their fellowship. I thank my contacts at Canvas and the team for their openness, warm hospitality, and participation during my time there. I am also indebted to all construction workers who supported this research by participating and sharing their knowledge. I learned a lot from each one of you.

An important part of my Ph.D. was the privilege to teach at the School of Architecture. I thank Steve Lee, Dana Cupkova, Omar Khan, and Sarosh Anklesaria for the opportunity to teach at the M.Arch program along with my teaching partners, Eddy Man Kim, Laura Garofalo, and Jonathan Kline. I am also grateful to my advisor Daniel Cardoso Llach for the opportunity to teach at the MSCD program. I learned a lot from you, Daragh Byrne, and later from Vernelle Noel as my teaching partners.

Solitude has perhaps been my biggest challenge in Ph.D. I appreciate my colleagues in the program, Jinmo Rhee, Manuel Rodriguez Ladron de Guevara, Pedro Veloso, Jingyang Liu, and Ardavan Bidgoli for easing it. I thank Ardavan for his early guidance, Jingyang for our long chats, Pedro (and Kate) for sharing their food, drinks, and conversation, and Manuel for his partnership in Pittsburgh and New York night life expeditions. I thank Jinmo for his lifelong friendship.

Last but not least, my heartfelt thanks go to my family and a special group of friends that I continue to grow up together. Thank you so much for being with me or on the other side of the phone during my past years of grind. Thank you for giving me strength by reminding me where I have come from and who I am truly surrounded by in life. Thank you for bringing a smile to my face each and every day. I am so lucky to have you all.

Table of Contents

1. Introduction 1

1.1 “Construction Work” in Robotics 3
1.2 Investigating Technology Design in Robotics 5
1.3 Questions 10
1.4 Sites and Methods 10
1.4.1 Overview 10
1.4.2 Site: Carnegie Mellon University Robotics Institute, Pittsburgh, Pennsylvania 13
Introduction and Relevance 13
Data Collection 15
Data Analysis 16
1.4.3 Site: The Rethinking Automation in Construction Project, Pittsburgh, Pennsylvania 16
Introduction and Relevance 16
Data Collection 17
Data Analysis 18
1.4.4 Site: Canvas Construction Robotics, San Francisco, California 20
Introduction and Relevance 20
Data Collection 21
Data Analysis 21
1.5 Contributions 22
1.6 Limitations 24
1.7 General Outline 25

2. Decomposing Construction Work and Reimagining It With Robots: CMU (1979–1989) 27

2.1 Outline 28
2.2 The Road to the 1980s 29
2.2.1 Robotic Arms and Computer Control 29
2.2.2 The Launch of the Robotics Institute 35
2.3 From Hard Automation to Programmable Automation: Robotics at the Institute 38
2.3.1 High-Level Languages, Reprogrammability, and the High Hopes for Manufacturing 38
2.3.2 Machines Operating Machines 41
2.4 Situating Robotics in Construction: Robotics at the Department of Civil Engineering 44
2.4.1 “Mobile Work Object” Versus “Stationary Work Object” 44
2.4.2 Perception, Navigation, and the High Hopes for Construction 45
2.4.3 Same Technology, Similar Discourse, Different Work 48
2.4.4 Decomposing Construction Work for Robot Development 55
A Taxonomy of Construction Work: Examining Warszawski and Sangrey’s Approach 56
A Modular Robot Design Framework: Examining Skibniewski’s Approach 58
2.5 Summary 61

3. Studying Construction Work Towards Robot Development: The ReAC Project 64

3.1 Outline 65
3.2 An Ethnographic Study of Construction for Robot Development 65
3.2.1 The New Scaife Hall Site 65
3.2.2 Descriptions of Building Tasks 68
Cleaning and Housekeeping 69
Fireproofing 72
Rebar Tying 76
Carpentry Formwork 79
3.3 Key Insights 83
3.3.1 Way of Working: Zones and Interactions 83
3.3.2 Qualities of Work: Grounding the Generic Descriptions 88
3.3.3 Worker Perspectives on Robots: Mixed Feelings 89

4. Reimagining the Role of Robots: The ReAC Project 93

4.1 Outline 94
4.2 An Ethnographic Study of Robotic System Development 94
4.2.1 Mill 19 94
4.2.2 Ethnography Drawing Robotics and Construction Closer 96
Bringing the Site Into the Robotics Lab: Prototype A – Data Viewer 97
On Site With Multiple Tasks in Mind: Collective Imagination and Conceptual Guides 97
A Brief Description of System Development Work During Ethnography 101
Robots for Ancillary Activities: Practical Guides, Use Cases, and a Scenario 103
From Site to System Development: Practical Inputs 114
4.3 Summary 122

5. Construction Work Informing a System: Canvas 124

5.1 Outline 125
5.2 An Ethnographic Study of Robotic System Development: Part 1 126
5.2.1 Canvas 126
5.2.2 Knowledge at the Core: Drywall Finishing 127
What Is Drywall Finishing, and Why Is It Needed? 128
Drywall Finishing Workflow and Tools 129
Levels of Drywall Finishing 133
Division of Labor and Skills 134
5.2.3 From Conventional to Robotic Drywall Finishing Workflow: Basil 136
A Brief Description of Basil 136
A Productive Swing Between Steps, Quality, and Technology 137
Sense-Think-Act + Evaluate: Work-Sharing Between Finishers and the Machine 140

6. System Reconfiguring Construction Work: Canvas 149

6.1 Outline 150
6.2 An Ethnographic Study of Robotic System Development: Part 2 150
6.2.1 On and Off Site: Reconfiguring Drywall Finishing Work 150
Tracing Worker–Robot Interactions: Prototype B – Workflow Visualization Tool 150
User Interface Next to Pan and Knife: Rearranging the Drywall Finishing Skill Set 152
Mapping Robot Development Work: Prototype C – Network Graph Visualization Tool 162
Labor at Canvas: New Participants, New Roles, New Knowledge 164
6.3 Summary 173

7. Discussion 175

7.1 Outline 176
7.2 Between Past and Current Conceptions of Construction Work 177
7.2.1 From High-Level Understandings to Situated Engagements With Construction Work 177
7.2.2 Machine Learning for Support: An Alternative to Single-Task Robot Approach 179
7.3 Between Past Aspirations and the Current Reconfiguration 180
7.3.1 From Visions of Automation to Robot Support and Worker–Robot Work-Sharing 180
7.3.2 Productivity as a Key and a Barrier 183
7.3.3 Emerging Technical Labor in Construction 184
7.4 Final Reflections 186
7.4.1 Reflections on Technology Design in Construction Robotics 187
The Situated Nature of Tasks 187
Skills, Deskilling, and Upskilling 188
The Fitness of Robots and Sites 189
Roles for Robots 190
7.4.2 Reflections on Pedagogy in Architecture, Construction, and Robotics 191
7.4.3 Reflections on the Configuration of Ethnography 193
7.4.4 A Final Note 194

Bibliography 195

Appendix 224

A- Key Definitions: What Is Work, What Is a Robot, What Is Robotics? 225
B- Extended Descriptions of Study Planning and Data Collection 227
Site: Carnegie Mellon University Robotics Institute, Pittsburgh, Pennsylvania 227
Data Collection 227
Site: The Rethinking Automation in Construction Project, Pittsburgh, Pennsylvania 230
Planning and On-Site Improvisation 230

Data Collection 231
Site: Canvas Construction Robotics, San Francisco, California 233
Planning and On-Site Improvisation 233
Data Collection 234
C- Descriptions of Building Tasks (continuing from Chapter 3) 236
Mechanical Electrical Plumbing (MEP) Ceiling Line Installations 236
Framing 243
Drywall Installation 248
D- Garbage Delivery and Concurrent Cleaning Use Cases (from Chapter 4) 252
E- MSD-RAS Seminar Course: Key Concepts and Questions Explored 253

Glossary

Key Definitions

- **Construction robot:** On-site construction robots described in this dissertation fall closer to Rodney Brooks's (2020) following "robot" definition: "Some sort of device which has sensors, which senses the world, does some sort of computation, decides on an action, and then does that action based on the sensor input, which makes some change out in the world outside its body."

- **Construction robotics:** Construction robotics is part of robotics that intersects architecture and civil engineering that focuses on the design, development, and applications of robots for construction.

- **Construction work:** The body of physical and mental activities regularly and professionally performed on materials or with instruments to execute construction tasks. Carrying out these activities requires training to learn the essential skills and techniques as well as practical experience to potentially master them over time.

- **Robot: "**An autonomous machine capable of sensing its environment, carrying out computations to make decisions, and performing actions in the real world" (IEEE, 2020). These technologies are interchangeably called as "robotic technologies," "robotic systems," and "robots" in this dissertation.

- **Robotics:** The scientific and technological research field that deals with the design, development, and application of robotic technologies in universities or industry. Rather than a branch of artificial intelligence (AI), robotics is defined as an inter-discipline between mechanical engineering, electrical engineering, and computer science in this dissertation.

- **Work:** Physical or mental activities performed on materials or with instruments and professionally as part of an occupation for pay (Schmidt, 2011; Ryle, 1951; Google-Oxford Languages Dictionary, 2024). Work and labor are used interchangeably in this dissertation.

See Appendix A for detailed descriptions.

Method Definitions

- **Archival research** is about basing historical research and writing on extracted primary-source documentary evidence such as oral accounts, documents, visuals, and artifacts (King, 2012).

- **Close reading** is a technique in humanities for systematic observation and interpretation of facts and details about a text (Kain, 1998). The technique is also carried into "software studies" with the aim of critically examining software systems as well as their social and cultural effects (Cardoso Llach at al., 2021). Close reading in this dissertation focused on technical and visual descriptions of robotic technologies together with perspectives, discourses, and future projections weaved around them in publications or other written archival materials.

- **Ethnography** consists in researcher's participation in a cultural context, often called as fieldwork, to investigate and describe particular phenomena by studying the actions, interactions, languages of actors along with physical or digital artifacts in that context (Hammersley & Atkinson, 1983; Bloor & Wood, 2006; Merriam, 2009). Participant observation, observation, interviews are the classic methods to collect documentary evidence —field notes, audio-visual recordings, oral accounts, and artifacts— in an ethnographic fieldwork.

- **Interview** is a data collection method that involves learning about a subject from participants' views, perceptions, or experiences. Interviews can be less structured to explore participants' views as well as semi-structured or structured to draw specific information from them (Brinkmann, 2008).

- **Observation** is a data collection method that involves observing a setting to gain an understanding of a situation or process. Observation in ethnography can be performed as an active participant (participant observation) or more of an observer than an active participant within the setting (observation) (McKechnie, 2008).

- **Oral history** is a method of historical information collection in the form of "memories and personal commentaries through recorded interviews" (Ritchie, 1994).

Chapter 1

Introduction

Advancements in robotics have led to the exploration of arms, mobile robots, drones, and other semi-autonomous technologies in building construction. In the United States, forecasts of an increase in revenue and the deployment of these machines in construction permeate the popular media.[1] Meanwhile, research and development efforts exploring their potential for on- and off-site operations continue within university labs and companies.[2] The software and hardware systems, models, and user interfaces that constitute these machines are everyday work artifacts, situated at the center of various professional groups concerned with their design. These include researchers and engineers collaborating to develop these systems as well as the laborers performing construction work via and around them.[3] The systems shaped by researchers and engineers inscribe transformations in the work of laborers, while the work of laborers is embedded within the system development processes of researchers and engineers.

Although a considerable portion of previous research is dedicated to the development of robotic systems, how the work of laborers has been engaged within those development processes and the resultant systems' role in reshaping that work have been underexplored. This dissertation offers a critical examination of "construction work" in robotics and the sociotechnical changes robotic technologies bring into construction work. This examination is important to understand and develop a critical view of how construction work has been understood and analyzed in robotics, and how robotic technologies have been discussed, developed, and integrated into construction accordingly. It is needed to initiate an early understanding of the slow, emerging reconfiguration of work around robotic systems. This critical view and understanding serve to inform both technology design and building practices and pedagogies.

The dissertation combines historical, ethnographic and computational methods, and engages three main sites where robotic technologies have evolved and continue to be explored in the United States: Carnegie Mellon University's Robotics Institute (CMU-RI, or the Robotics Institute) and Department of Civil Engineering (DCE) in the 1980s, a construction robotics research project named Rethinking Automation in Construction (ReAC) at Carnegie Mellon University (CMU), and a San Francisco–based construction robotics company named Canvas. Historical research is employed to examine a moment in the 1980s, where a group of civil engineers at the intersection of CMU-RI and the DCE were building the field of construction robotics in the United States, technically and discursively. Ethnographic research extended the examination into the field's present. The two

1 For instance, A study by research firm Tractica (2019) (cited in *Business Wire* (2019)) has estimated that the global construction robotics revenues will reach from $22.7 million in 2018 to $226 million in 2025 fueled by the growing interest in the architecture engineering construction (AEC) industry. The study also estimated more than seven thousand robotic devices and equipment assisting various construction tasks, building and demolishing structures, and handling construction finishing work on construction sites by the same year. Parallel to this growth globally, there is labor shortage in construction in the United States. A survey by the Associated General Contractors of America (2022) has revealed ninety-one percent of contractors experience difficulties in finding workers.

2 In the United States, Robotics & Intelligent Construction Automation Laboratory (RICAL) at Georgia Tech, Construction Automation, Robotics & Visualization Laboratory (CARV Lab) at Auburn University, Construction Automation and Robotics Laboratory (CARL) at North Carolina University, Automation & Robotics in Construction and Design (ARCADE) at Virginia Tech are some examples focusing on the development of fundamental technologies, while Digital Architecture Research and Technologies (DART) Laboratory at University of Michigan, Autonomous Manufacturing Lab (AML) at University of Pennsylvania, Adel Design Research (ADR) at Princeton University, Design Fabrication (DFab) Laboratory at Carnegie Mellon University are among examples focusing on applied fabrication research. In the industry, one may look at the emerging group of construction robotics companies such as Advanced Construction Robotics, Built Robotics, Canvas, Construction Robotics, Dusty Robotics, Nextera Robotics, SafeAI, and Toggle.

3 Collaborative work via and around computational systems has been an existing notion and an object of study. This notion has been the foundation of computer-supported cooperative work (CSCW), a research field that stemmed from computer science in the 1980s, and focuses on the research of computational systems for collaborative work and collaborative work through computational systems. See Paul Wilson (1991) and Jonathan Grudin (1994) for the origins of CSCW as well as David Pinelle and Carl Gutwin (2000), Jacques Wainer and Claudia Barsottini (2007), and James Wallace and coauthors (2017) for reviews of the field.

contemporary practices were selected to participate in their processes engaging both construction and technology design work.

By exploring technology designers' conceptions of construction work in these sites, the dissertation identifies a shift from high-level descriptions and classifications of building tasks concerned with their general character and commonalities at CMU to close, situated engagements with workers and knowledge of tasks aimed at understanding their specificities at ReAC and Canvas. Through examining how past aspirations to reconfigure construction at CMU unfold at present, the dissertation traces the emergence of alternative approaches seeking robot support and worker–robot work-sharing for sensible integration of machines at ReAC and Canvas, next to bold visions of automation and productivity increase prevailing since the 1980s. Further, it reveals that rearrangement of skills along with an emerging robotics labor intertwining with a building trade constitute the current reconfiguration around Canvas's state-of-the-art machine, rather than the labor displacement projected in the 1980s at CMU.

Drawing from these findings, the dissertation develops three arguments. The cases demonstrate that technology designers' conceptions of work, alongside technical artifacts, configure their analyses and shape system development in construction robotics. Relatedly, participation of workers and designers in each other's knowledge worlds frame systems with alternative roles (compared to that of archetypal single-task robots) and systems which show a greater potential to sensibly fit into existing workflows on sites. Lastly, automation progresses faster in designers' discourse than it does via machines, evidenced by the discrepancy between the past aspirations and the current realization in the field. With these, the dissertation contributes a historically and socially grounded understanding of how we conceive robots and of the related reconfiguration of work, encouraging new considerations, technologies, and pedagogies across computational design, human–computer interaction, robotics, and architecture.

1.1 "Construction Work" in Robotics

What do I mean by "construction work" in robotics, and how is construction work embedded in system development processes? Work consists of skilled activities that may be as basic as decision-making, movement, or sensing. Oftentimes, the goal in robotics is to emulate a selection of these, which are necessary to perform labor, with a system. A general approach to achieving that goal has been to represent the enactment of these activities and the environments in which they take place with data and models. An example is Ci-Jyun Liang and coauthors' (2020) learning from demonstration (LfD) experiment, which extracted state-action data—the pose of a unit tile and next action—from an expert tile installation demo to train an industrial arm with reinforcement learning in simulation.[4] The representation of skills and context was also fundamental in developing past robotic technologies. For Instance, Mark Cutkosky and Paul Wright (1986) at Carnegie Mellon University Robotics Institute studied and modeled the grip mechanics of machine operators to develop a

4 (1) Some other examples demonstrating this general approach (involving the representation of skills and context) in robotics include: Deborah Clever and collaborators' (2018) use of motion capture to represent human gait and produce a model-based description of gait towards gait generation in robots, Kai Liu and collaborators' (2018) analysis of arm-reaching movements to extract movement primitives as a basis for the mechanical design of an exoskeletal arm, or Marcin Andrychowicz and collaborators (2020) reinforcement learning-based system that determines the next in-hand manipulation action of a five-finger gripper based on its initial state and an image-based pose estimation of the held object. (2) Based on robotics researchers Aude Billard and Daniel Grollman's (2013) entry in Scholarpedia, "robot learning from demonstration (LfD)" (that Liang and coauthors (2020) used) is about end-users showing robots the desired behavior to be performed as opposed to decomposing and programming the behavior. In computing terms, LfD, as Brenna Argall and colleagues (2009) describe, is a technique that involves learning a policy (a mapping between the world state and actions) from demonstrations (example state-to-action mappings). See Argall and colleagues' (2009) survey of LfD.

robotic gripper combining power and dexterity.[5] Any of these skilled actions in the examples is essentially decontextualized from the labor that it constitutes.[6] For instance, a generic arm motion action appears as a unit that constitutes hammering a nail once we recontextualize it in that labor.

Thus, in this dissertation, I approach construction work not just as an end goal, a form of performance, that robotic systems are developed to accomplish. It is not just a sociotechnical operation that is affected by the end products of robotics. Construction work has also been a subject in the making of robots, more specifically, a source of knowledge that engineers and researchers have sought to understand in developing and informing software and hardware, models, and user interfaces.

If construction work is a subject in system development, how is it studied or analyzed by engineers and researchers who are not its inherent laborers? Science and technology studies (STS) scholar Madeleine Akrich (1992) points out the "hypotheses" that technology designers build about the "entities that make up the world" in which they deploy "their [technical] objects." Part of their job consists in "inscribing this vision of (or prediction about) the world in the technical content of the new object," says Akrich (1992). In the context of construction robotics, the entity that is continuously re-envisioned is construction work. Engineers and researchers in robotics build conceptions of this entity, and these conceptions are reflected in the ways they study, analyze, and understand it as well as inscribed into the systems they build. The first line of investigation in this dissertation focuses on these conceptions in the history and current practice of construction robotics: What conceptions of construction work are formed and guide these analyses and, accordingly, the conception and development of on-site robotic systems?

Alongside technical developments, the field of robotics has been producing discourses that contain aspirations to reconfigure work with machines. Sociotechnical concerns regarding work, including enhancing productivity, saving labor costs, increasing labor safety, overcoming the limitations of labor, and compensating for labor shortages, have been the default motivations of technical research and systems towards automation. For instance, Miroslaw Skibniewski and Mani Golparvar-Fard (2015) highlight the "waste in coordinating labor and equipment" and "managing, moving, and installing material" in the construction industry and underscore the industry's productivity decline for more than three decades. Yue Pan and Limao Zhang's (2021) review is another example calling attention to the scale of the construction industry and the financial value that it may gain from an increase in productivity via AI- and robot-based automation.[7] Similar discourses about work have also been part of robotics in the past. In the 1980s, engineers Abraham Warszawski and Dwight

5 Several other examples (that involve representation of skills and context) exist in the history of robotics. For example, as a basis for complex object recognition by robots, Gerald Agin (1972) at the Stanford Artificial Intelligence Laboratory (SAIL) developed a model that represents curved objects with an axis and a series of cross-section planes. For the lab's well-known mobile robot, called the "Stanford Cart," Hans Moravec (1980) built a rudimentary obstacle detection system, which compares sample regions in recorded images of the environment and computes three-dimensional positions of objects to be incorporated into the system's world model. At the Massachusetts Institute of Technology (MIT) Artificial Intelligence Laboratory, Terry Winograd (1971, 1973) developed an early language model based on representations of our knowledge about syntax, semantics, and deductive reasoning, which enabled a robot arm to interact with users in natural language and make inferences during pick-and-place tasks in simulation.

6 Here, the term "decontextualized" is borrowed from anthropologist of technology Diana Forsythe (2001). Forsythe (2001) discusses AI scientists' "construction of their experts' work" (toward developing expert systems) as "decontextualized," meaning that they see "little distinction between observable practice and verbal descriptions of that practice."

7 A similar example in robotics include Anthony Downs and collaborators' (2021) article evaluating the state of industrial robots for cutting labor costs and facilitating cost-competitive "low-volume/high-mix production" among small and medium-sized manufacturers in the United States. Another one is P. M. M. Santos and collaborators' (2021) article discussing automation as a prerequisite for reviving textile manufacturing in developed countries, which typically outsourced this work due to domestic labor shortages or costs. In construction robotics, Nathan Melenbrink and coauthors' (2020) review of on-site robotic systems points out the advantages of these technologies for improving productivity and labor safety in a list of tasks.

Sangrey (1985) and William Whittaker (1986) made a case for using robots in construction by foregrounding the industry's low productivity despite its large labor force.[8]

Consequently, aside from attending to construction work as a knowledge source in the making of on-site robotic systems, I also approach it as a sociotechnical operation that is reconfigured in discourse during the development of these systems and, more recently, in practice through their deployment. This means that the activities that constitute work or the arrangement of actors and devices that facilitate their enactment, such as workers, instruments, interactions, and organizational structures, are not just technically but also discursively reproduced by these figures.

These discursive reproductions are themselves part of what Akrich (1992) describes as "inscribing… in the technical content of the new object," which, in this case, is a robotic system. Thus, in Akrich's (1992) terms, I simultaneously explore "the world inscribed" in robotic systems and "the world described" by their deployment in construction practice in this dissertation. The second line of investigation I follow concerns the reconfigurations of construction work conceived in the past or emerging from current construction practice around robots: What kind of reconfigurations are aspired to and attached to the conception and development of on-site robotic systems? What kind of reconfigurations begin to emerge around these systems in current construction practice?

1.2 Investigating Technology Design in Robotics

Despite their presence in construction robotics, conceptions of building construction work, the ways these conceptions guide the development of robotic systems, and the reconfiguration of construction work envisioned or beginning to emerge via these systems have not been explored in depth.

Construction robotics research has mainly focused on toolmaking, material fabrication, and automation efforts.[9] Toolmaking efforts concentrate on prototyping tooling and software geared towards improving the usability of robots for building tasks. Recent examples include a flexible five-fingered end effector by Charlotte Firth and collaborators (2020), a real-time robot control software by Jose Luis Garcia (2019), and a deep learning–based motion control by Gabriella Rossi and Paul Nicholas (2019).[10] Material fabrication efforts explore the relationship between tool, craft, and material through data, as shown by Selen Ercan Jenny and collaborators' (2020) robotic plaster-spraying experiments, or materials' structural and performative attributes, like Aleksandra Anna Apolinarska and collaborators' (2021) free-form timber frame construction.[11] Automation efforts concerned with moving repetitive labor-intensive tasks into robotic workflows include Juhyeon Kim

8 (1) I further discuss these key figures along with their points on productivity in construction in Chapter 2. (2) The history of robotics includes other cases. For instance, George Devol, the lead inventor of the first industrial robot arm in the United States, suggested that robots could improve autoworkers' conditions by taking over their risky and repetitive tasks while also decreasing costs (Ballard et al., 2012). In the 1970s, MIT researcher Marvin Minsky (1972) critiqued the traditional assembly line as the "major consumer of unskilled labor" and proposed robot-based "individualized mass production" for improving the United States' global trade position. In the 1980s, a cost comparison of workers and robots on industrial tasks, such as operating machines or inspecting pieces, was featured in the engineer Joseph Engelberger's (1980) famous book *Robotics in Practice*. Additionally, some earlier publications, such as those by Robert Ayres and Steve Miller (1981), A.W.J. Chisholm (1987), and Abraham Warszawski (1985), also show the field's concerns about the implications of robots in manufacturing and construction labor and the economy.

9 Based on a high-level review of related publications in *Cumulative Index about Publications in Computer Aided Architectural Design (CuminCAD)* database, *Construction Robotics*, and *Automation in Construction* journals.

10 Other notable examples include Nathan Melenbrink and collaborators' (2017) autonomous strut-climbing robot, Andrew Moorman and collaborators' (2016) pipeline for real-time context-sensing of effectors, and Johannes Braumann and Sigrid Brell-Cokcan's (2015) KUKA Prc, a parametric robot control and simulation software.

11 Other notable robotic fabrication projects include Maria Yablonina and collaborators' (2021) filament fabrication experiments exploring new possibilities in formal aesthetics and reconfiguration of space with task-specific robots, and Joshua Bard and collaborators' (2018) robotic concrete surface printing towards an improved thermodynamic performance.

and collaborators' (2022) 3D reconstruction method enabling scaffold construction monitoring via robot dogs and Aiyu Zhu and collaborators' (2021) object-tracking method for multi-robot housing assembly.

These studies are key to demonstrating robotic systems' potential for building small-scale structures or fabricating their parts. They are concerned with these systems as instruments of technical improvement or design rather than tools that entail shifts in labor and organization in building design and construction. Only a limited set of works in construction robotics deal with, alongside robots' enabling capabilities, the connections between these technologies and the innate elements that constitute construction work, such as craft, knowledge, and the enactment of human–machine interactions.[12] Among these works are Joshua Bard and collaborators' (2014, 2016) explorations reattaching plasterers' hand motion skills to robotic plastering workflows through motion tracking and Daniel Cardoso Llach and coauthors' (2017) study highlighting the tactile and embodied aspects of student-authored fabrication processes. Questions of technology design and work related to robotics, such as how robotic systems are developed and how they might affect construction work, have not been critically explored extensively by the construction robotics research community.

In parallel, a group of researchers in architecture and computational design have looked into questions of technology design and work in the context of building design and construction, but the focus of their investigations has been on computer-aided design (CAD) systems. For instance, Paolo Tombesi (2010) examined how information technology has increased the value of project management skills for architects and shifted their position in the organizational division of building design. Focusing on the collaborations in a US practice, Carrie Dossick and Gina Neff (2011) showed the building information modeling (BIM) software's limitations in supporting flexible problem-solving talk and tacit knowledge exchange between architects, engineers, and construction professionals. Combining archival research on the history of the technical development of CAD with participatory observation in a construction project in Abu Dhabi, Cardoso Llach (2015, 2019) confronted the historical projections that the software would facilitate a seamless integration between design and construction actors. Through an in-situ investigation, Yanni Loukissas (2012) explored how computer simulation practices have shaped new professional identities and how simulation software functions within the social system of knowledge in a global architecture and engineering practice.

These investigations have shown the intrinsic yet often hidden connections between technology design and work related to CAD software in building design and construction. Some of these have expanded the methodology in tracing these connections by incorporating historical and ethnographic research along with computational techniques for examining software and work. However, the knowledge produced by this strand of research has so far remained specific to CAD systems. It has provided valuable accounts of the conceptions of work in the making of these technologies, as in Cardoso Llach's (2015) work, or on how these technologies inscribe transformations of labor and organizations in building design and construction, like Tombesi's (2010) or Loukissas's (2012) research. The strand has not yet been extended to investigating the conceptions and reconfigurations of work pertaining to other technologies that are less mainstream and newly emerging in building design and construction, such as artificial intelligence (AI) and robotic systems.[13]

12 Based on the *CuminCAD*, *Construction Robotics*, and *Automation in Construction* review.
13 See other notable studies by Dominik Holzer (2009), Peggy Deamer (2012), Sarmad Mehrbod and coauthors (2019).

Drawing from the anthropology of technology as well as science and technology studies (STS), other fields including robotics/human–robot interaction (HRI), human–computer interaction (HCI), and computer-supported cooperative work (CSCW) have investigated the connections between technology design and work. While most robotics/HRI research focuses on technical development, a small group of studies at the intersection of anthropology, STS, and the field have critically examined the cultures of technology design in robotics. The roots of these studies lie in early STS research on the cultures of scientific and technology design work in laboratories.[14] A seminal example is philosopher Bruno Latour and sociologist Steve Woolgar's (1979) ethnographic study revealing how scientists' everyday activities condition the ways scientific knowledge is constructed at the Salk Institute for Biological Studies. Recent scholarly work in this lineage has produced critical reflections on the views, values, and assumptions as well as abstractions framing knowledge representation in robotics. For instance, humanities scholar Jennifer Rhee (2018) discussed how researchers' familiarity with a form of knowledge or labor and their implicit cultural views of it shape the way knowledge is represented in a system. Similarly, humanities scholars Neda Atanasoski and Kalindi Vora (2019) foregrounded the culturally biased and techno-utopian conceptions of humans and labor surrounding robots.

Closer to the inquiry in this dissertation, a subset of these studies has traced various conceptions of engineers and researchers in robotic system development through ethnographic engagements with labs and companies.[15] For example, anthropologist Bohkyung Chun and computer scientist Heather Knight's (2020) study analyzed the role of "anthropomorphic descriptives" in self-driving vehicle teams in relating to each other and to the systems they develop. Sociologist Arne Maibaum and coauthors' (2021) research on five projects across Europe and the United States demonstrated roboticists' decontextualization of elderly care practice and how this leads to systems fitting into care scenarios that are more contrived than actual ones. The study by HRI researcher Selma Šabanović (2007) involved robotics research institutes in the United States and Japan and traced "cultural conceptions of social interaction" in developing and evaluating social robots. Anthropologist Catherine Hasse and collaborators' (2019) multi-sited study with roboticists, users, and other "affected stakeholders" in five application domains (including construction) in Europe revealed the discrepancy between the makers' robot design descriptions and robots in use.[16] So far, research on conceptions in robotics/HRI has not been fully extended to construction robotics and its history in the United States.[17] Specifically, conceptions of construction work in relation to the development of on-site machines in the United States have not been explored.

---

14 See STS scholar Karin Knorr-Cetina's (1983a, 1983b) reviews of these early laboratory studies. See Wiebe Bijker and coauthors' (eds.) (1987) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology* as well as Bijker and John Law's (eds.) (1992) *Shaping Technology/Building Society: Studies in Sociotechnical Change* for a collection of critical studies focusing on scientific and technology design work. See also a collection of digital STS works, that extend these classic studies into contemporary digital technologies, practices, and contexts, in Janet Vertesi and David Ribes's (eds.) (2019) *Digital STS: A Field Guide for Science & Technology Studies*.

15 I found this small set of ethnographic studies (focusing on technology designers' conceptions in robotic system development) through reviews conducted in (A) *ACM Digital Library*, (B) *Paladyn Journal of Behavioral Robotics*, (C) *American Society of Civil Engineers (ASCE) Library*, and (D) *Automation in Construction* journal. The review in *ACM Digital Library* focused on *ACM/IEEE International Conference on Human-Robot Interaction* proceedings and *ACM Transactions of Human-Robot Interaction (THRI*) journal that yielded the most relevant studies. The review in *ASCE Library* focused on *Journal of Computing in Civil Engineering*, *Journal of Construction Engineering and Managemen*t, and *Journal of Management in Engineering* that returned results. Lastly, the search in *Automation in Construction* did not yield any relevant works. See Eunjeong Cheon and Norman Makoto Su (2016, 2018), Omar Mubin and coauthors (2018), Hee Rin Lee and coauthors (2019), Niels Christian Mossfeldt Nickelsen (2019), Jessica Sorenson (2019) as some other studies that I did not expand in the text.

16 The difference of this study from the one in this dissertation is twofold: (1) The focus is the European context, (2) although the technologies the researchers studied existed, the cases in which they traced the work of designers and stakeholders around these technologies were "not real cases" due to "strict non-disclosure agreements" between the researchers and the companies (Hasse et al., 2019).

17 Based on the review of ethnographic studies in (A), (B), (C), (D) in endnote #15. Additionally, the existing historical research are surveys of technologies, such as robot arms in the United States and Europe or construction robots in Japan, as opposed to

In a similar vein but from a technical standpoint, computer scientists and roboticists in robotics/HRI have also produced critical reflections, underlining the lack of contextual approaches in developing and evaluating robotic systems. For instance, Rodney Brooks (1991) critiqued the way researchers conceptualize systems by "abstracting away most of the details to form a simple description" of a situation. Later, Brooks (2018) also criticized the distance between user studies in labs or on crowdsourcing platforms without robots and real-world contexts. Along the same lines, Maja Matarić (2018) pointed out the difficulty of formalizing "real-world problems" and representing them faithfully in the design of experiment–type system evaluations. How lab- and crowdsourcing-based evaluations of human–robot interaction overlook actual users and more complex use contexts was also emphasized by Kerstin Dautenhahn (2018). Offering similar critiques, Malte Jung and Pamela Hinds (2018) invited the field to engage with work, domestic, and public spaces, while Sara Kiesler and Michael Goodrich (2018) called for ethnographic studies in addition to lab-based studies to understand users' interaction with systems. Building on these critiques of the gap between systems and real-world contexts, a close look into how engineers and researchers engage with the work practice for which they develop technologies could be informative for robotics. Yet, this question remains underexplored in the specific context of construction robotics. How construction work is situated within robotics and how robotic technologies are situated within construction work in the United States have not been examined.

Similarly, although technical development has been the main focus in robotics/HRI, a small group of studies across HCI, CSCW, and HRI have attended to the effect of robotic systems on work.[18] In HCI and CSCW, research on computational systems and work has been a longstanding preoccupation to produce knowledge and technology design implications.[19] It took some years for these fields to extend their studies, which initially leaned heavily on design-of-experiment methods, to the real world.[20] Researchers began to study the social and material practices around computational systems, as in Lucy Suchman's (1993) study in an airline operations room and Christian Heath and Paul Luff's (2000) research in a newsroom, a subway control room, and an architectural office.[21] More recently, a number of studies in HCI and CSCW have concentrated on robotic systems. However, a large portion of these are user perception or experience studies intended to produce implications for design and—like in the fields' earlier years.[22] For example, Marina Escobar-Planas and collaborators' (2022) study analyzed children's perceptions of social robots relative to robot behavior, and Simon Coghlan and collaborators (2021) elicited elders' perceptions through videos of companion robots and interviews.[23]

works dealing with the histories of the field across a variety of sites, including actors, processes, technologies, and technologies' effects, as previously noted.

18 I found this group of studies (focusing on the effect of robotic systems on work) based on the review of (A), (B), (C), (D) in endnote #15, additionally, through a review of (W) *ACM Conference on Computer Supported Cooperative Work (CSCW)*, (X) *Proceedings of the ACM on Human-Computer Interaction (PACMHCI/CSCW)* journal, (Y) *Computer Supported Cooperative Work (CSCW)* journal, and (Z) *ACM Conference on Human Factors in Computing Systems (CHI)*.

19 See the history of this research within HCI/ CSCW in Jonathan Grudin's (1994) "Computer-supported Cooperative Work: History and Focus," Graham Button and Wes Sharrock's (2009) *Studies of Work and the Workplace in HCI: Concepts and Techniques*, and Kjeld Schmidt and Liam Bannon's (2013) "Constructing CSCW: The First Quarter Century."

20 It is important to note that, other than the early ethnographies of science and technology in STS previously noted, theoretical constructs such as "situated action" by Lucy Suchman (1987) and "cooperative work as design" by Susanne Bødker and collaborators (1988) had major influence on the field to undertake in-situ research of cultures and technological systems of work. See Paul Dourish and Graham Button, (1998), Randall and coauthors (2005) for reflections on ethnography in HCI/CSCW, and Jeanette Blomberg and Helena Karasti (2013) for a reflection and review of ethnographic works.

21 As part of this early strand, see also Genevieve Filippi and Jacques Theureau (1993), Richard Harper (2000). For early works that carried this approach into technology design, see John Hughes and collaborators (1994), Jeanette Blomberg and collaborators (1996), Lucy Suchman and collaborators (1999). For recent examples situated at the intersection of STS and HCI/CSCW, see Morgan Ames (2019), Janet Vertesi (2019), Angele Christin (2020).

22 Based on the review of (W), (X), (Y), (Z) in endnote #18.

23 See Prajna Bhat and Yuhang Zhao (2022), Shuai Yuan and coauthors (2022), Richard Paluch and Claudia Müller (2022) as some other studies among others (focusing on user perception or experience through) that I did not include in the text.

Next to these user perception and experience studies, a small number of works have investigated technological reconfigurations around robots in work settings.[24] These include Hannah Pelikan and coauthors' (2018) ethnographic study in US hospitals, which showed that working with the Da Vinci surgical robot rearranged the placement of members inside the operating room and the distribution of information among them. Another example is Amy Cheatle and coauthors' (2019) study, which uncovered how working with the same robot shifts operators' eyesight to sight through the system and reduces the importance of tactile skills in surgery. Similar research was conducted in manufacturing by Eunjeong Cheon and colleagues (2022a), who documented how cell-style production with collaborative arms generated new worker interactions within each shift and across shifts, along with new worker–robot support staff and worker–management interactions in a Danish facility. Therefore, the canonical question of how computational technologies affect collaborative work in HCI and CSCW has recently begun to be investigated around robotic systems. Yet, these investigations have so far focused on healthcare and manufacturing work practices and particular types of robots used for this work, as in the examples above.[25] The effects of on-site robots on construction work have not been researched within these fields.

Similar to HCI and CSCW, HRI is dominated by studies focusing on users' perceptions of and experience with robots.[26] Recent examples explore these themes with various social robots; for instance, Sara Nielsen and collaborators' (2023) study used digital ethnography on YouTube videos, and Anastasia Ostrowski and collaborators (2022) employed interviews with elderly people together with their robot log data.[27] Additionally, HRI features a number of studies aimed at finding robot use cases in various work practices, such as Leon Bodenhagen and collaborators' (2019) ethnographic observations at a Danish hospital, Katie Winkle and collaborators' (2018) focus groups and interviews with therapists in the United Kingdom, or Elizabeth Cha and collaborators' (2015) interviews with inhabitants in the United States about their domestic work organization.[28]

Conversely, studies that focus on the technological reconfigurations of work around robots are rare in HRI.[29] One recent example is Jamie Wallace's (2021) interviews with manufacturing workers in Denmark, which showed that "learning," "values," and "skills" dynamically emerged with the integration of collaborative robots and role changes from operating to servicing around these tools. Another of these studies is Marie Opdal Ulset and coauthors' (2022) interviews with elderly care workers, which revealed that the use of the social robot Pa-sonal Robotto (PARO) added new labor rather than functioning as an automation tool for elderly care in a New Zealand retirement home. Also focusing on PARO, Wan-Ling Chang and Selma Šabanović's (2015a) ethnographic study captured how the robot's use in a US nursing home affected workers' "daily routines" with the

24 Based on the review of (W), (X), (Y), (Z) in endnote #18. The reason why these studies are not many may possibly the current limited presence of robots in work settings as well as limited researcher access to work settings that have robots. See Pieter Duysburgh and coauthors (2014), Eunjeong Cheon and coauthors (2022b), Paul Hatch and coauthors (2023) as other studies (focusing on the technological reconfigurations of work around robots) that I did not include in the text.

25 Based on the review of (W), (X), (Y), (Z) in endnote #18. CSCW includes a small group of studies that have attended to building design and construction work—see Kjeld Schmidt and Ina Wagner (2004), Melanie Tory and coauthors (2008), Lars Rune Christensen (2013). However, these studies have not necessarily focused on robotic systems.

26 Based on the review of (A), (B), (C), (D) in endnote #15.

27 See Iolanda Leite and coauthors (2012), Hee Rin Lee and Selma Šabanović (2014), Antonia Meissner and coauthors (2020), Matthew Beane (2020), Hannah Pelikan and coauthors (2020), Sofia Thunberg and Tom Ziemke (2020), Franziska Fischer and coauthors (2023) as some other examples among many (focusing on user perception or experience) that I did not include in the text.

28 See Caroline Pantofaru and coauthors (2012), Wan-Ling Chang and Selma Šabanović (2013), Kory Kraft (2016), Tommaso Colombino and coauthors (2020) as some other notable studies—aimed at finding use cases in work practices—that I did not include in the text.

29 Based on the review of (A), (B), (C), (D) in endnote #15.

residents.[30] Lastly, a couple of early investigations by Jodi Forlizzi and Carl DiSalvo (2006) and Ja-Young Sung and coauthors (2008) documented the ways the Roomba cleaning robot influences housekeeping practices in the United States. Like in HCI and CSCW, healthcare and manufacturing work and robots have been the primary foci of these studies, with the addition of domestic work practices.[31] Technological reconfigurations of construction work around on-site robots have not yet been studied in the field.

The two lines of investigation in this dissertation, which focus on the conceptions and reconfigurations of construction work in robotics, aim to address these gaps. Through a historical and ethnographic approach, they engage in two types of work in the selected sites as objects of study. The first is construction work, which I frame as both a source of knowledge in robotic system development and a sociotechnical operation that is potentially reconfigured by robotic systems. The second type of work that the investigation engages in is robotic system development itself. I study this work in history and participate in it in current practice to investigate technology designers' understandings of construction work and their aspirations to reconfigure it. Lastly, the investigation focuses on the robotic workflows along with the software and hardware that constitute machines to examine the current reconfiguration around these technologies.

1.3 Questions

The investigation in the dissertation is built around these questions:

- Q1A: How do technology designers conceptualize and analyze construction work to inform the development of robotic systems in history and current practice? How do these conceptions come into play in system development?
- Q1B: What are the linkages or shifts between past and contemporary conceptions of construction work and systems?
- Q2A: What kind of reconfigurations do robotic systems entail in building construction?
- Q2B: What are the alignments or shifts between past aspirations to reconfigure construction work and emerging reconfigurations around state-of-the-art systems?

1.4 Sites and Methods

1.4.1 Overview

To investigate the questions above, I conducted mixed-method research combining historical, ethnographic, and computational methods. The research engaged primary sites: CMU-RI and DCE in the 1980s, the ReAC project at CMU, and the San Francisco–based robotics company Canvas (Figure 1).

The first part of this mixed-method effort involved historical research, through which I explored question Q1A by focusing on the CMU-RI and the DCE (Figure 1). I picked these sites due to their historical significance in robotics and in the emergence of construction robotics in the United States. I focused on the 1980s because of the importance of this decade in the history of robotics. I explain the relevance of this site and the period further in the next section. My approach to studying the

30 See Alessandra Maria Sabelli and coauthors (2011), Wan-Ling Chang and Selma Šabanović (2015b), Lasse Blond (2019), Ned Barker and coauthors (2020) as some other notable examples (focusing on the technological reconfigurations of work around robots) that I did not include in the text.

31 Based on the review of (A), (B), (C), (D) in endnote #15.

history of robotics combined media archaeology and software studies perspectives with conventional historical research methods.[32] The data collection relied on conventional methods, including archival research to extract the broader histories of robotics, automation, and the sites or biographical details about key figures but also on oral history to incorporate first-hand accounts from a subset of these individuals (see Appendix B). In addition, I treated technical research papers, robotic systems that were the subjects of these papers, and documentation of these systems (such as software architecture diagrams) as historical evidence for some of the discussions in the dissertation. I also used close readings of software as a method for analyzing these systems (see Appendix B). This enabled me to balance the systems' social and technical aspects in my analysis.

The second part of the mixed-method effort consisted of ethnographic research, through which I investigated questions Q1A and Q2A on two contemporary sites: ReAC, in which I actively participated, and Canvas (Figure 1). These two were selected out of five potential sites in the United States. They were well-fitted for this research because together, they covered the early-stage development of a prototype and the late-stage development of a commercial-grade robot, robot deployment, and engagements with construction sites. This allowed me to look into the conceptions of construction work in system development and the effects of systems on construction work in parallel. Access and the duration of fieldwork were also important criteria in selecting these sites. Robotics is an expensive and challenging enterprise, which makes it difficult to locate a site with an ongoing project and a timeline permitting long ethnographic work. While ReAC provided me with an opportunity for active participation for two years, Canvas granted me access to their office for two and a half months in person and remote access to some of their resources for another five months. Approximately eight to ten months of two years were active fieldwork across construction sites, labs, robot test cells, and the online work platforms of ReAC and Canvas.

While the ReAC study was designed in two layers, namely, an ethnographic study of construction work informing the development of a support robot and an ethnography of the whole process, the Canvas study involved an ethnographic study of the company's development and deployment work around its drywall finishing robot. The ethnographic study of construction work within the framework of the ReAC project builds on the practice-based technology design lineage in HCI/CSCW; therefore, the study was geared for informing a system in the making.[33] Both the ethnographies of ReAC and the ethnography of Canvas concerned with approaches to technology design build on the ethnographies of science and technology that originated in STS and later infiltrated into HCI/CSCW

32 A concise description of media archaeology I have found is in the synopsis of Erkki Huhtamo and Jussi Parikka's (eds.) (2011) book *Media Archaeology: Approaches, Applications, and Implications*, which identifies it as "an archaeological approach" to studying "how media were written about, used, designed, preserved, and sometimes discarded." A similar compact description for software studies is provided by Wendy Hui Kyong Chun and collaborators (2022): an area of study that "uses and develops cultural, theoretical, and practice-oriented approaches to make critical, historical, and experimental accounts of (and interventions via) the objects and processes of software." See more in Matthew Fuller's (2008) introduction in *Software Studies: A Lexicon* along with Mark Marino's (2020) *Critical Code Studies*.

33 Lucy Suchman (2002b) identifies practice-based technology design as "[technology] design based in everyday practices," in which designers "locating themselves within the extended networks of social relations and forms of work that constitute technical systems." The aim is to inform the design of systems with an understanding of social and material aspects of work practices. The practice-based technology design lineage in HCI/CSCW goes back to Suchman's and Jeanette Blomberg's pioneering efforts at Xerox Palo Alto Research Center (PARC)—see Suchman (1983), Blomberg and collaborators (1996), Suchman and collaborators (1999). Some recent examples of practice-based work in robotics/HRI include William Juel and coauthors' (2020) development of an elderly care robot concept based on ethnographic and participatory design work at a Danish care center, Marcus Mast's (2014) development of assistive robot interfaces based on a set of qualitative studies with elderly people in Europe, and Hyunjin Kim and colleagues' (2007) development of a path planning approach for home cleaning robots based on a study of cleaning behaviors in Korea. As an addition to this relatively small body of work, ReAC's ethnographic study of construction work I conducted carries this approach into construction robotics.

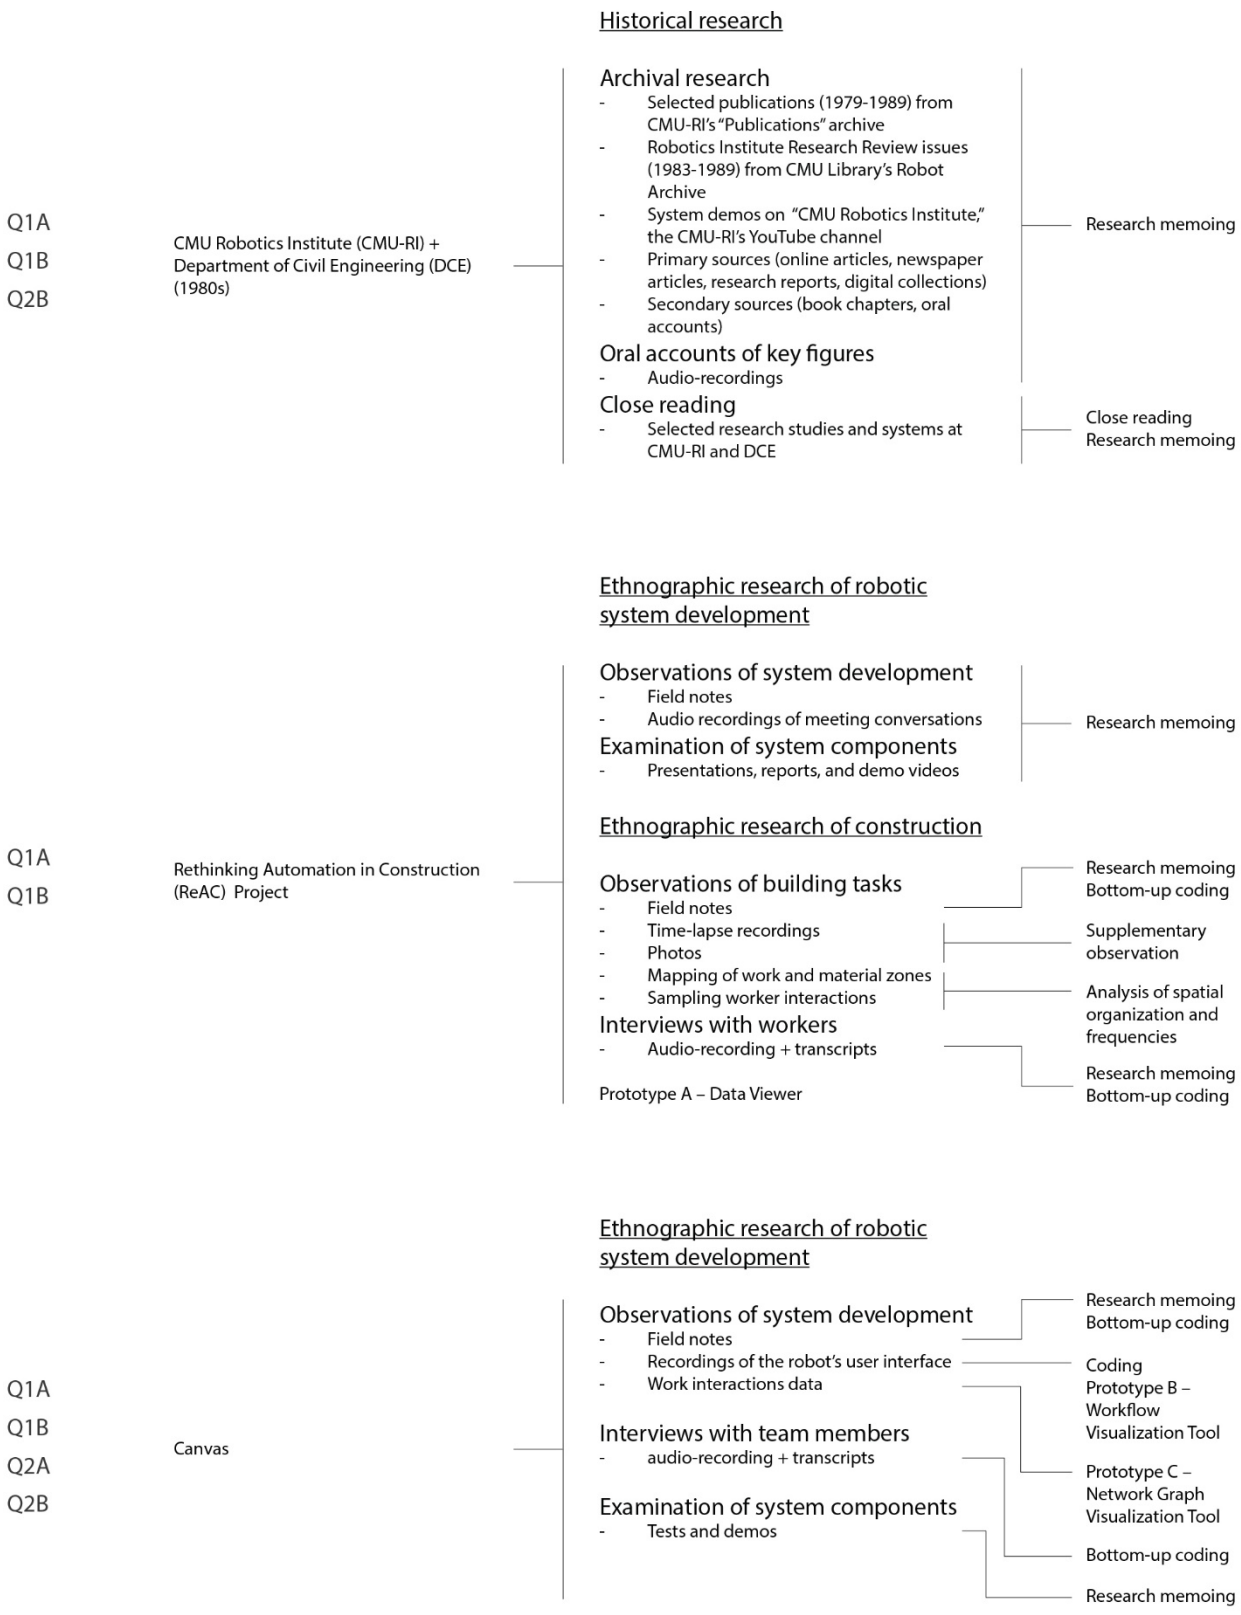


Figure 1. Questions, sites, and methods in the dissertation

and HRI to produce accounts of the cultures of technology design in computing.[34] Aside from the lineages they build on, both studies involve a "design (and computing) in ethnography" approach, informed by the recent ethnographic work in computational design and HCI/CSCW.[35] In this approach, I used computational techniques, in particular data science, network analysis, and visualization, and prototyped research tools in support of conventional field methods, including participant observation, observation, and interviews. Through these techniques and prototypes, I expanded the observational data collection from physical spaces hosting participants' in-person work interactions to digital spaces hosting their interactions with each other or with software. Further, I explored creative ways to represent field data.

Lastly, through the historical research focusing on CMU and the ethnographic research on ReAC and Canvas, I investigated questions Q1B and Q2B about linkages and shifts between construction robotics in the 1980s and today (Figure 1).

1.4.2 Site: Carnegie Mellon University Robotics Institute, Pittsburgh, Pennsylvania

Introduction and Relevance

The historical research for this dissertation focused on the CMU-RI and the DCE, whose members established a line of construction robotics work through a close relationship with the Robotics Institute. In terms of the historical period, I briefly explored the 1960s and 1970s to draw the broader context of robotics and labor in the United States into which the CMU-RI was born. Then, I closely looked at the 1980s, starting with the launch of the Robotics Institute in 1979 and covering the key figures and work in its first decade alongside those of the DCE.[36]

I selected the CMU-RI because of its significant role in the history of robotics in the United States. As the first dedicated institute in a US university, it pioneered the institutionalization of robotics research and education in the country.[37] A highly active center in the 1980s, it concentrated on basic and applied AI/robotics research for manufacturing and hazardous work. The researchers at the Robotics Institute produced several systems historically significant to manufacturing automation as well as some of the country's early mobile robots significant to autonomous navigation research in the 1980s. Some of these became objects of study in this dissertation. As a forefront engineering knowledge hub in the United States, the Robotics Institute also played a part in improving the country's industrial productivity through sponsorships by the Westinghouse Electric Corporation and the Department of Defense's research arm, the Office of Naval Research (ONR).[38] The work under this agenda ranged from evaluating the implications of industrial robots for labor to developing a

34 (1) In that sense, the second ethnography in ReAC and the ethnography in Canvas are versions of ethnographies computer scientist Paul Dourish and anthropologist Genevieve Bell (2011) characterize as free from the pressure of "requirements-gathering" for a technology product, and instead informing with "implications for how we go about design in the first place, what it might do, how it can be carried out, and what import it holds." In Dourish and Bell's (2011) words, "engagements with topics, people, and field sites" in these two studies "are used to understand phenomena of significance to design," and in both cases, about technology design itself. (2) A group of works that I reviewed in this chapter have extended this lineage in STS and HCI/CSCW into studying technology design in robotics—see Selma Šabanović (2007), Eunjeong Cheon and Norman Makoto Su (2016, 2018), Catherine Hasse and collaborators (2019), Bohkyung Chun and Heather Knight (2020).

35 I explained this approach in Section 1.5.

36 The launch year of the Robotics Institute is based on the brief information section titled "History of the Robotics Institute," published in the Carnegie Mellon University Robotics Institute website (2022). See the source in Bibliography.

37 The fact that the Robotics Institute is the first dedicated one in a United States university is based on two sources: (1) a description section titled "Robotics Program" in Carnegie Mellon University's "Undergraduate Catalog," published in the University's (2024a) website, and (2) a brief description section titled "Cruise Controlled," published in the University's (2024b) website under "About." See these sources in Bibliography.

38 As I describe further in Chapter 2, Westinghouse Electric Corporation and the Office of Naval Research were the initial funders of the Institute (Carnegie Mellon University Robotics Institute, 2022).

flexible manufacturing cell for a Westinghouse turbine plant and researching autonomous navigation systems for ONR towards potential military applications.

The Robotics Institute was also a crossroads for researchers from CMU's engineering departments. The relationship between the Institute and the DCE in the 1980s made CMU a prominent center of mobile robotics research in the United States and one of the few sites in the world where construction robotics emerged as a subfield of robotics. The DCE played a pioneering role in organizing the global research community in this emerging field and, consequently, in its institutionalization as an academic research field. As I describe in the next chapter, the DCE famously organized and hosted the first two symposiums on construction robotics, named the "International Symposium on Automation and Robotics in Construction" (ISARC), in 1984 and 1985, which were the seeds for the foundation of the International Association for Automation and Robotics in Construction (IAARC) in 1990 (IAARC, 2021). In the 1980s, the DCE, together with the Robotics Institute, also spearheaded the production of construction robot prototypes in the United States. Due to its significant position and production in the early years of construction robotics, I considered the DCE worthy of attention for the questions in this research.

The reason for putting the magnifying glass on the 1980s was the decade's importance to robotics in three ways. First, a list of advancements that came about in the 1970s matured in the 1980s and multiplied robots' value for the manufacturing industry. The key one was computer control and the proliferation of high-level robot programming languages, which improved the ability to reprogram an arm-based manipulator for different tasks with a computer and easier-to-use languages, as opposed to reprogramming its operating system.[39] Second, with computer control, high-level languages, and improved sensing, the 1980s were the decade when robots gained more attention in the industry and became emblematic vehicles for automation and the United States' global industrial competitiveness, particularly against Japan.[40] Lastly, the 1980s were also the time when on-site construction robotics started to emerge as an organized research field and community due to the leading efforts of CMU's DCE.[41] This stood as an alternative to the corporate-led efforts in late 1970s Japan that produced a set of single-task robots (Bock & Linner, 2016).[42] This community,

39 Other advancements in the 1980s were the developments in vision and force sensing capabilities along with AI techniques in control systems, all geared towards handling more complex tasks with robots (Ayres & Miller, 1981; Gasparetto & Scalera, 2019).

40 For instance, a past report by Ayres and Miller (1981) titled "The Impact of Industrial Robots" (that I further discuss in Chapter 2) reported a survey with the members of the Robot Institute of America, which revealed "to keep up with the Japanese" as one of the motivations for robots in the industry.

41 A series of sources mark the 1980s as the beginning of construction robotics. According to researchers Thomas Bock and Silke Langenberg's (2014) brief historical overview, the first construction automation efforts came from Japanese manufacturers, such as Sekisui, Toyota, and Panasonic, and the initial results were off-site prefabrication facilities (structuring manual labor in an assembly line setting) as early as the 1960s. As Bock's overview with Langenberg (2014) and his later account with Thomas Linner (2016) state, the interest in Japan shifted from off-site automation to research and development of on-site robots, namely single-task robots, in mid to late 1970s, initially led by the contractor Shimizu, then followed by the Ministry of Trade and Industry, other contractors, and universities. The initial efforts to develop these robots started in the late 1970s, but the publications by Tomohiro Ueno and colleagues (1986) as well as Miroslaw Skibniewski and Jeffrey Russell (1989) show the very first functioning prototypes were only ready and introduced to construction sites in the early 1980s and onwards. A similar effort but in a modest scale started at CMU's DCE in the United States around those years as Borja Garcia de Soto and Miroslaw Skibniewski (2020) note, and I detail in this dissertation. These efforts initiated in Japan and spread to the United states and Europe began to institutionalize around the International Symposia on Automation and Robotics in Construction (ISARC) by the leading efforts of the DCE also in those years (IAARC, 2021). With all these facts, researchers consider the idea of single-task robots as the root of construction robotics, and the early1980s, when these robots are introduced to sites in Japan and the research gets organized in academic circles with the leading efforts of CMU, as the beginning of the field—see Carl Haas and coauthors (1995); Kamel Saidi and coauthors (2016); Borja Garcia de Soto and Miroslaw Skibniewski (2020).

42 A single-task robot in construction robotics nomenclature is what we generally refer to when using the term "construction robot." Thomas Bock and Thomas Linner (2016) define single-task robots as "systems that support workers on the construction site in executing one specific construction process or task or by completely substituting the physical activity of human workers necessary to perform this one process or task." Though the definition of these technologies may slightly vary in other sources. For instance, Bock and Linner (2016) used the phrase "systems that support workers … in executing one specific construction process or task" next to "completely substituting the physical activity." In a slightly different way. Borja Garcia de Soto and

which was distributed across Japan and Europe as well as the United States, where it was led by CMU and MIT, began to produce and publish more prototypes in this decade (Skibniewski & Russell, 1989).[43] With all these advances and prototypes unfolding new possibilities for manufacturing and construction industries up to and throughout the 1980s, discourses on labor became more enmeshed into the discourses on technology.

By looking into the civil engineers' discourses, analyses of work, and robotic systems in this context, I investigated their conceptions of construction work prevalent during the initial years of construction robotics at CMU. I explored their aspirations to reconfigure work via systems while situating robotics in construction.

Data Collection

At the center of historical research was a two-part investigation of three major archives: the CMU-RI's "Publications" archive, an online repository that includes published works from the Robotics Institute since 1979; the CMU Library's Robot Archive, a mixed-media archive featuring annual research reports titled the *Robotics Institute Research Review* published between 1983 and 1994; and "CMU Robotics Institute," the CMU-RI's official YouTube channel, which contains video clips of robotic system demos and researcher interviews from the 1980s and 1990s.

The first part entailed site data collection to draw CMU's institutional context. It consisted of a high-level review of robotics publications between 1979 and 1989 from the Publications archive, the *Robotics Institute Research Review* issues between 1983 and 1989 from the Robot Archive, and the information narrated in the interviews and system demos on the YouTube channel (Figure 2). These sources, especially the *Robotics Institute Research Review* issues, along with online articles, newspapers, oral accounts helped trace the details of the Robotics Institute's launch, sponsors, and motivations (Figure 2). Reviewing the publications along with these other sources allowed me to

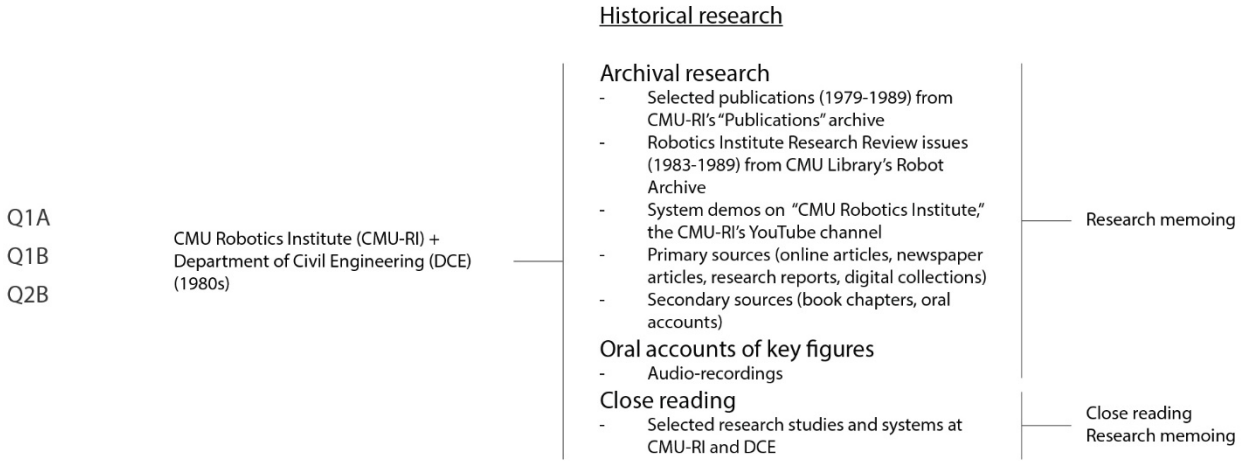


Figure 2. Data collection and analysis (CMU-RI and DCE)

Miroslaw Skibniewski (2020) describe these robots as "designed to imitate skilled labor," and "the main purpose" of these technologies as "to alleviate skilled labor for those jobs considered "3k" or "kitsui" (difficult), "kitanai" (dirty), and "kiken" (dangerous)." See a survey of these robots in Bock and Linner (2016).

43 See a list of construction robot prototypes developed within and outside the United States by 1989 in Skibniewski and Russell (1989).

identify and document the key figures active at the intersection of the Robotics Institute and the DCE (See Appendix B). To position this local institutional context, I also referred to primary sources such as research reports, newspapers, digital collections, and secondary historical sources that informed me about the broader context of robotics and labor in the United States (Figure 2).

The second part involved a more extensive survey of publications on the work produced in CMU's institutional context. The review covered the CMU-RI and DCE publications between 1979 and 1989 as well as the 1980s *Robotics Institute Research Review* issues and allowed me to extract a set of key research studies and systems (of historical significance and relevance to the investigation) for close reading and examination (See Appendix B). To obtain further details about CMU's institutional context and projects, I collected oral accounts from a subset of the key figures previously listed (Figure 2) (See Appendix B). Other than these, I relied on two existing oral accounts from the founding director of the CMU-RI, which provided valuable details about the Institute's launch.[44]

Data Analysis

The data collection accumulated a variety of materials, including robotics publications, research reports, system demo videos, oral accounts, online articles, newspaper articles, and secondary historical sources about the site, key figures, and robotic systems. Their analysis had two main objectives. One was to delineate the broader context of robotics and labor in the 1960s and 1970s, which created the conditions leading to the Robotics Institute's launch, along with the local context of the Institute and the DCE, which gave rise to the emergence of construction robotics at CMU in the early 1980s. The other objective was to identify conceptions of work and aspirations to reconfigure it in the early days of construction robotics at the DCE. For analysis, I used research memoing on the written sources and oral accounts (Figure 2). I combined it with a close reading of a series of robotic arms, which were the focus of robotics in the 1960s and 1970s, arm-based manufacturing systems and mobile robots developed at CMU in the 1980s, and proposed construction robotic systems and research studies analyzing construction work developed at the DCE in the 1980s (Figure 2).

The analysis helped build the brief account of the changes in robotics and labor related to key advancements in robotic arms in the United States presented in Section 2.2.1 and Section 2.3.1. It also enabled the extraction of details about the Robotics Institute's launch, including sponsors and underlying motivations, which I describe in Section 2.2.2, and showed the emergence of a discourse about reconfiguring manufacturing work around systems at the Institute, as described in Section 2.3.2. The analysis also facilitated a discussion of the effect of early mobile robots on the emergence of construction robotics as a field at the DCE in Section 2.4.2 and Section 2.4.3. It helped to describe this emergence around key figures at the Department and trace in their work aspirations to reconfigure construction work, which I detail in Section 2.4.3. Lastly, the analysis allowed me to trace DCE engineers' conceptions of construction work in their discourse, analyses, and systems, as reported in Section 2.4.3 and Section 2.4.4.

1.4.3 Site: The Rethinking Automation in Construction Project, Pittsburgh, Pennsylvania

Introduction and Relevance

ReAC is an interdisciplinary research group based in the Computational Design Laboratory at the School of Architecture at CMU. The group focuses on developing technical frameworks for adaptive

44 These two oral accounts from the founding director of the Robotics Institute are obtained by John Markoff (2018), a historian at the Computer History Museum, and Hansen Hsu (2021), a curator at the Computer History Museum Software History Center.

robots to explore human-centered, multi-robot, or human–robot interactive workflows in building construction. The group consists of researchers with diverse multidisciplinary backgrounds who combine knowledge of robotics, machine learning, mechanical engineering, computational design, construction, and qualitative methods. I was an active member of the group for two years from September 2021 to December 2023.

The ReAC project that became part of the ethnographic research in this dissertation took place between September 2021 and May 2023 in collaboration with a large regional construction company. The project aimed to explore a practice-based approach to designing a robotic system for construction and, through this system, demonstrate a robot-in-the-loop scenario alternative to the prevalent full automation or human-in-the-loop scenarios in construction robotics. To do so, the project incorporated ethnography to closely study construction work and engage with workers and, consequently, inform the conception and development of a prototype robotic system for supporting a specific task. As the member leading this effort between the site and the lab, I designed and carried out this study on an educational building construction site managed by the project collaborator. Guided by the data and insights acquired from the site, the group ultimately developed a reinforcement learning (RL)–based navigation system and deployed it on a mobile robot to support carpentry formwork by performing tool/hardware delivery. Finally, we evaluated the system with construction worker participants.

As a rare case in construction robotics, this project conjoining ethnographic research on construction work and system development appeared as an appropriate site for investigating the questions in this dissertation. It promised both an investigation and active participation in the development of a system by handling its ethnographic component and conducting an ethnography of both the ethnographic and technical components that distinctively constitute the system's development. Through observation and participation in this contemporary robotics project, I explored our conceptions of construction work and automation partly shaped by the ethnographic study, and how the study and these conceptions guided the way we conceived, developed, and situated a robotic system within a specific task.

Data Collection

Designed in two layers, the ethnographic research in the ReAC project consisted of an ethnographic study of construction work and an ethnographic study of robotic system development, which the first study aided (Figure 3). The planning of the first study involved identifying data collection methods, communicating these methods to the general contractor and its subcontractors as participants, scheduling on-site data collection, and obtaining the required permissions. I initially identified four main data collection methods: on-site observations of building tasks supplemented by time-lapse recordings, photos, and site mappings; interviews with selected workers, questionnaires to capture worker interactions, and a technology design workshop between workers and the ReAC group. Due to extremely tight job schedules, it was not possible to reserve workers' time for more than one interview, so the questionnaires and design workshop were later dropped from the list of methods.

Like for this first study, the planning for the second study focusing on the robotic system development work included the identification of data collection methods, scheduling, obtaining the required permissions, and communicating these methods to my ReAC colleagues as participants. I initially identified five main data collection methods: participant observations of system development activities, including weekly meetings, discussions, and hands-on work, the examination of the system modules and hardware that would constitute the prototype robot, interviews with selected

colleagues, experience sampling via questionnaires to extract the group's collaborative work interactions, and collecting samples from group conversations via audio recording.[45] Eventually, colleagues' views expressed naturally in those work moments and captured in these notes and audio recordings replaced the formal interviews. The questionnaires and collection of conversation samples were scheduled to match the duration of the first study, whereas the observations, examination of the system, and interviews were scheduled over a longer timeline.

Data Analysis

The analysis of the data from the first ethnographic study aimed at providing our group with an understanding of construction work and site conditions, identifying use cases and building a scenario for robot support, and concurrently informing the development of the prototype system. It consisted of a mix of individual and collective activities.

In parallel with each week's data collection, I analyzed the field notes through research memoing and bottom-up coding, which included initial/open coding (Figure 3).[46] Through this, I drew out detailed descriptions of the floor we selected to study (named Level B) and the observed tasks under "site description," "activities," "materials and tools," and "interactions" codes. I supported this analysis with quick remote observation via the time-lapse recordings, which helped fill the gaps in in-person observations (Figure 3). I also examined Level B's daily spatial organization into work and material storage zones based on the sketches and photos and the interaction frequencies of workers within and across the trades I recorded in sample observations (Figure 3). This first round of analysis produced empirically rich descriptions of the observed tasks. These descriptions presented in Section 3.2.2 informed the group about the site conditions in which the tasks were performed, the activities that constituted each task, the coordination of these activities among workers, and the movement of workers, materials, and tools.

The task descriptions were shared with the members in biweekly site information sessions, in which I presented one or a set of them. To support the sharing of the descriptions, I developed an interactive data viewer that combines and displays the various types of data collected from Level B, such as the time-lapse recordings of tasks, photos of worker activities and tools, drawings of work and material zones, and network visualizations of worker interactions (Figure 3). This research tool, which is described in detail in Section 4.2.2, is built around the idea of design and computing supporting ethnography in a practice-based system development process. By effectively bringing the construction site into the robotics lab, the prototype facilitated and enriched our discussions during the site information sessions, in which we collectively evaluated the opportunities and limitations of robot support in each task and identified the use cases discussed in Section 4.2.2. The key insights from Level B described in Section 3.3 were also drawn up and noted during the site information sessions around this prototype.

Interviews occurred after the observations; similar to observational data analysis, interview data analysis proceeded concurrently with interview data collection. With each interview, I performed a first round of analysis on the audio recording via research memoing (Figure 3). This allowed me to verify the information obtained from the observations. It also enabled me to extract initial points from those recordings that would enrich the task descriptions. For further analysis, I transcribed the recordings by using Zoom's audio transcription feature. Along with the interview notes, I used the

45 See Reed Larson and Mihaly Csikszentmihalyi (2014) for an overview of experience sampling method.
46 See Johnny Saldana (2013) for detailed descriptions of qualitative coding methods.

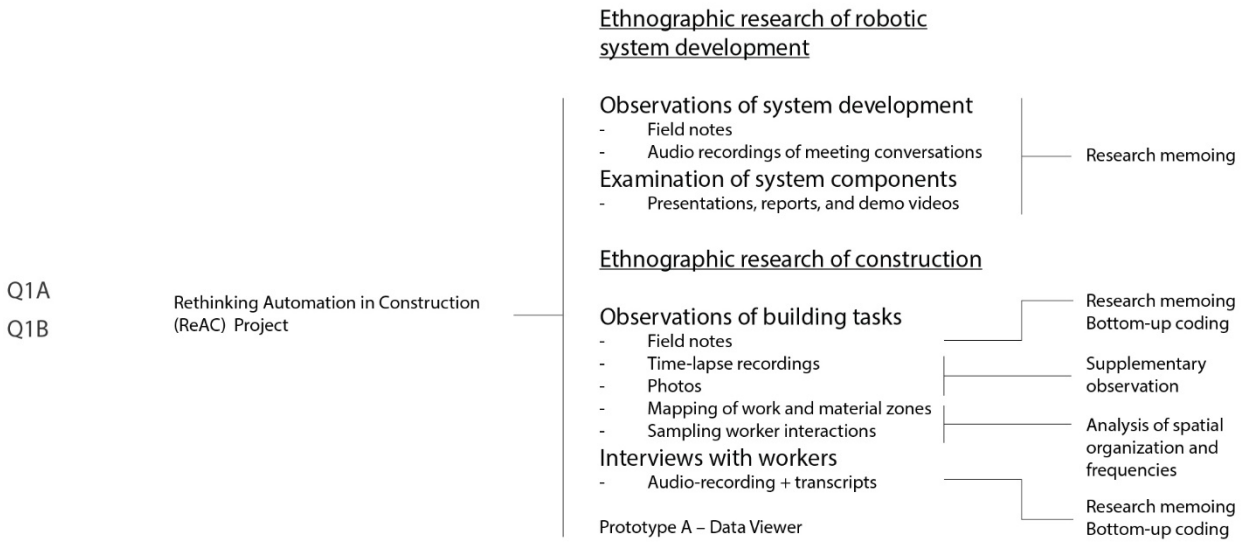


Figure 3. Data collection and analysis (ReAC)

same four codes I extracted from the field notes (site description, activities, materials and tools, and interactions) to index the data in these transcripts. This was followed by bottom-up coding, which involved initial/open coding to elicit the workers' knowledge of their tasks and helped derive new codes identifying the qualitative character of their actions, including "intellectual or physical," "individual or coordinative," "planned or improvisational," and "mechanical," "mobile," and "dexterous" (Figure 3). The bottom-up coding also involved value coding to capture the workers' perspectives on robots for construction, which is discussed in Section 3.3.3.

During site sessions, we evaluated the observed tasks in terms of robot support opportunities and noted preliminary use cases in each task, as described in more detail in Section 4.2.2. Then, together with my colleague leading the system development, I evaluated the feasibility of these cases based on the group's technical capabilities and the cases' practical value assessed initially by us and later by the workers during the interviews. Tool and hardware delivery, an existing supporting activity in carpentry formwork, emerged from this evaluation as a credible form of support that a robot could perform. A further analysis, which included documentation of the work zone in which formwork took place on Level B and coding part of its time-lapse recording towards producing a worker action log, was conducted. The log helped us validate if the support role we conceived of for the prototype in the making might be useful. I developed this use case into a detailed robot support scenario, described in Section 4.2.2, which guided parts of the implementation, testing, and evaluation of the system.

The analysis of the data from the second ethnographic study aimed at building a reflective account that examines the group's conceptions of construction work and automation shaped by the ethnographic study and how these conceptions guided the group's system development process. This account allowed me to reflect on the project's approach to analyzing construction work, conceiving a system, and positioning it in construction practice, in conjunction with the approaches followed in the historical case and the second contemporary case (Canvas) in the dissertation.

Research memoing was the primary method I used to analyze the field notes and various project materials, such as presentations, reports, and demo videos (Figure 3). While memoing, I extracted the significant points that would later help in delineating the conception of the final system based on

the ethnographic study on Level B. These points would also aid me in describing the system components, focusing on the data and insights from the study informing their implementation and testing. These descriptions are featured in Section 4.2.2. The aim of the questionnaires was to derive a network analysis of the group's collaboration to supplement the ethnographic descriptions of the process. For this purpose, I conceived a prototype research tool that would process this data and visualize the members' interactions through network graphs. However, I later considered this analysis not valuable for this study and instead built and used the network graph tool during the Canvas study for a similar analysis with a different type of data. Analyzing the transcriptions of our meetings and site information sessions helped us empirically discuss a key part of the system development, that is, our conversations in which construction and robotics knowledge converged. I used research memoing to analyze these conversations, and ultimately trace the ways the knowledge of building tasks we gained from Level B informed our conception of the system (Figure 3). The details of this analysis incorporated into the narrative in Section 4.2.2.

1.4.4 Site: Canvas Construction Robotics, San Francisco, California

Introduction and Relevance

Started in 2017, Canvas is a construction robotics practice, which has developed a drywall finishing robot called "Basil." Basil is a semi-autonomous system designed for spraying and sanding drywalls to deliver level 5 and level 4 finishing, which are standard in the construction industry. The body of the system is made of a human lift modified into a mobile base with sensors and a collaborative robot arm, which works with an airless sprayer and vacuum sander tools mounted on the base. The human lift also carries a compressor for the sprayer and sander as well as the vacuum of the sander and a computer running the system's software. The controller of the system consists of an off-the-shelf tablet that works as the user interface and a gamepad for steering the base. At the time I engaged with Canvas in 2022, Basil had already been used on several construction sites in the Bay Area.

Working around Basil is an interdisciplinary team that includes business professionals, mechanical, robotics, software, and electrical engineers, electrical technicians, fabricators, former construction professionals, and drywall finishers who work as robot operators. The engineering team I focused on in my study is the largest team, comprising a major portion of these professionals, including the in-house finishers. The team is made up of sub-teams: the Software Team, which is in charge of the late-stage development and maintenance of the system's software components; the Hardware Team, which works on the system's mechanical, hardware, and sensing components; the Process Team, which is tasked with the development and evaluation of the robotic workflow; and the Operations Team, which takes care of job site deployment, training, support, and maintenance operations. The Engineering Team also includes a group of managers leading the entire team and the sub-teams.

In the United States, a limited number of construction robotics companies are potentially available for researchers who would like to study them. Among these, Canvas is one of the few to possess an advanced-level product that is deployed on sites, a well-sized interdisciplinary engineering team, and an open mindset towards hosting a researcher to shadow them. By giving me access, Canvas offered me the opportunity to investigate the questions in this dissertation by focusing on the late-stage development of their system and collaborative work interactions within their engineering team. The company permitted me to observe their meetings, tests, and a job site deployment and provided access to a set of computational data to extend these observations through custom-built research

tools. They also allowed me to examine their system. In addition, the members of the engineering team, including the in-house finishers, volunteered to share their knowledge in interviews. Through these engagements, I investigated the conceptions of drywall finishing work that come into play in the development of Basil, how these conceptions shape the development of the system, and how the system reconfigures drywall finishing work.

Data Collection

The Canvas study involved data collection in two broad categories: (1) collection of data related to Basil's development and (2) collection of data related to the broader work organization around the system. It relied on observations, computational data collection, and interviews, all concentrating on the company's engineering sub-teams (Figure 4). The methods I initially proposed to my site contacts at Canvas were observations of system development activities, including daily meetings, tests, and construction job site deployment(s), interviews with selected team members who conduct these activities, experience sampling via questionnaires to document their collaborative work interactions, and sampling of team conversations from online communication applications. Additionally, an examination of Basil's software and hardware components as well as sampling recordings of the robot's user interface and robot system logs were included to expand the boundaries of data collection. At the site contacts' request, I removed sampling of team conversations. Finally, because the number of participants made it time consuming to collect questionnaires, the meeting calendar app—in which the company included me to follow the events—became an improvised data source for analyzing work interactions.[47]

Data Analysis

Similar to the ReAC study, the data analysis aimed at creating a reflective account tracing the conceptions of drywall finishing work at Canvas and how these conceptions shape the development of Basil. This account enabled me to examine the engineering team's approach to analyzing work, conceiving the system, and positioning the system in construction practice together with the approaches in the historical case and the first contemporary case (ReAC) in this dissertation. The account also analyzes the ways Basil, as a state-of-the-art system, reconfigures drywall finishing work. This allowed me to reflect on this current reconfiguration in practice as an example, next to past aspirations to reconfigure construction work with robots at the beginning of construction robotics at CMU.

Through field notes, I captured a wide range of information during the observation of meetings, tests, and the job site deployment (Figure 4). I employed descriptive coding to categorize these notes daily or weekly into the previously listed themes of the study—system development and the broader work organization around the system. Further, I added short reflective memos to these notes throughout the fieldwork (Figure 4). After the fieldwork, I used initial/open coding to extract the salient points from the notes (Figure 4). These points aided me in delineating the key observations during the system evaluation test and job site deployment, which are presented in Sections 5.2.3 and 6.2.1, respectively. They also helped describe, first, the overall system and the translation of conventional drywall finishing to robotic drywall finishing workflow in Section 5.2.3 and, second, the ways the system reconfigures drywall finishing work in terms of skill set and organization in Section 6.2.1.

47 The company uses Google Calendar, a commonly used scheduling app by Google. The app's user interface includes a calendar, populated with meeting tabs that the user adds and invites other attendants, or alternatively gets invited to. Each tab holds the time, attendants, location, and other meeting information. More information on Google Calendar app is widely available online.

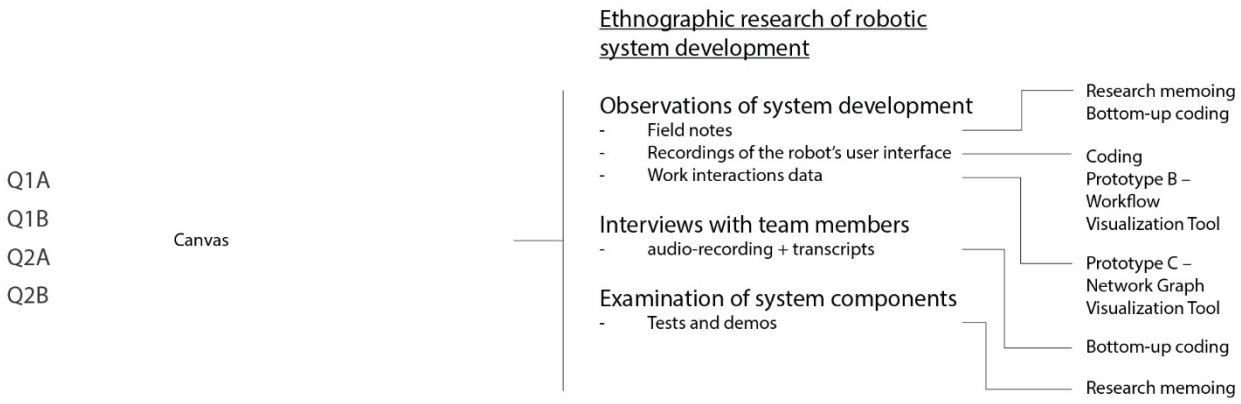


Figure 4. Data collection and analysis (Canvas)

The interview data provided many essential details for constructing the reflective account (Figure 4). The interviews were structured based on a set of broad themes, including the participants' knowledge of their work, their perspectives on their work and on robots for construction, computational aspects of the system, details of the robotic drywall finishing workflow, and the role of construction knowledge in system development. After organizing the data according to these themes, I used initial/open coding for analysis (Figure 4). I performed the coding by listening to the audio recordings, excerpted the coded verbatims from the transcripts, and noted the important points. This analysis helped elicit the in-house finishers' knowledge of the work, which I present in Section 5.2.2, and trace how the engineers conceptualized finishing work based on this knowledge and how they used this knowledge while developing the system, which I discuss in Section 5.2.3. The analysis also helped me build the two key discussions, on the new skills and new technical labor emerging around the system, presented in Section 6.2.1. To analyze part of the interviews, I also used value coding to uncover the participants' perspectives on their work and robots. These perspectives are examined in Section 6.2.1.

Lastly, I used two custom-built research tools to analyze the computational data. This analysis supported the discussions on the new skills and new technical labor in Section 6.2.1. To analyze the user interface (UI) screenshots, I designed a qualitative coding scheme that consists of the standard Basil UI actions along with the worker and robot actions observed during the job site deployment. Based on this scheme, I created a detailed time-stamped coding of the robotic drywall workflow on the job site (Figure 4). I built a workflow visualization tool, described in Section 6.2.1, to process this textual data and visualize worker–robot interactions (Figure 4). Part of the discussion on new skills in Section 6.2.1 was supplemented by this analysis and visualizations. Once I created this data, the system logs remained much less granular and redundant, so I did not analyze them. To analyze and visualize the work interaction data that I created based on the meeting calendar app, I developed a network graph visualization tool, presented in Section 6.2.1 (Figure 4). I relied on this analysis and visualizations alongside the observations to familiarize myself with the organization, plan my engagements with work activities and participants, describe the collaborative work around Basil, and build the discussion on the labor at Canvas in Section 6.2.1.

1.5 Contributions

Situated between computational design, construction robotics, and architecture, this dissertation contributes to the intersection of these fields. The historical and ethnographic studies in this dissertation add a line of sociotechnical research on the development of robotic technologies in the

United States to the existing toolmaking, material fabrication, and automation efforts in these fields.[48] The final discussion offers a snapshot of how construction work has been engaged in the field and how past aspirations to reconfigure work unfold at present, which may potentially inform the research, technology design, and pedagogical practices. The ethnographic studies document an instance of an ongoing change in building construction instruments and labor, by tracing the conceptions of work, the rearrangement of skills, and the interdisciplinary technical labor in robotics, which increasingly participate in construction. In doing so, they extend the previous research that investigated similar changes in the past around software to construction robotic systems.[49]

The dissertation's other contribution to computational design, construction robotics, and architecture are two sets of reflections on technology design and pedagogy featured in the final chapter. Reflections on technology design focus on approaches and methods for pursuing a sensible technology-work integration in construction robotics as a response to four key issues revealed in the ethnographic studies: the situated nature of tasks; skills, deskilling, and upskilling; the fitness of robots and sites; and roles for robots. Reflections on pedagogy underscores the importance of incorporating concepts and methods for studying, understanding, and articulating new questions about technology design, technologies, and technologies' effects on work into the pedagogies of these fields. Relatedly, a salient contribution of this dissertation is a seminar course developed for the Master of Science in Design: Robotics and Autonomous Systems (MSD-RAS) program at the University of Pennsylvania (UPenn). Titled "Histories-Reconfigurations-Representations-Reflections," it offers an introduction to the intellectual and technical origins of the intersection of CAD/AI/robotic systems with building design and construction as well as the current configuration of this intersection through readings, writings, and conversations. The seminar aims to educate future scholars, researchers, or toolmakers who can operate at the crossroads of design, computing, and the scholarly fields and who may investigate new questions similar to those in this dissertation.

Further, the dissertation adds to the growing research on technology design, robots, and work across HCI, CSCW, and robotics/HRI. Despite growing interest in critical histories of technologies in these fields, histories of construction robotics in the United States remain underexplored.[50] The historical research in this dissertation, which offers a retrospective view of past conceptions of work in one of construction robotics' prominent sites in the United States, is a first step towards filling that gap. In parallel, ethnographies of technology design infiltrating HRI from HCI and CSCW have begun to investigate the influence of engineers' and researchers' various conceptions in robotic system development.[51] The two ethnographic studies that trace the conceptions of work in two system development processes expand this emerging scope, which has mainly focused on manufacturing and healthcare work, to construction work in the United States. HCI, CSCW, and HRI have also

48 Example works reviewed in Section 1.2.

49 See some of these studies in Section 1.2.

50 Critical histories of technologies in HCI/CSCW is in formative stage—see the discussion by Robert Soden and coauthors (2021) foregrounding history as an "under-utilized area of focus." This strand of research is therefore limited to a small group of studies—see the brief survey of existing work by Soden and coauthors (2019, 2023)—and has not yet attended to the histories related to robotics and robotic technologies. Aside from these existing works of critical history, HCI/CSCW include other reviews documenting the emblematic technologies or the evolution of the field—see Jonathan Grudin (1990, 1994, 2017), Brad Myers, (1998), Ronald Baecker (2008), Schmidt and Bannon (2013). On the other side, robotics/HRI lacks research focusing on the histories of the field and technologies, which is pointed out by researchers Alessandro Gasparetto and Lorenzo Scalera (2019). The existing historical accounts by Johanna Wallén (2008), Bruno Siciliano and Oussama Khatib (2016), Irati Zamalloa and coauthors (2017), and Alessandro Gasparetto and Lorenzo Scalera (2019) are brief surveys offering the evolution of technologies, rather than detailed historical or sociotechnical accounts engaging critical questions related to technology design. Accounts reporting on the history of construction robotics by Leslie Cousineau and Nobuyasu Miura (1998), Thomas Bock (2007), Thomas Bock and Silke Langenberg (2014), Thomas Bock and Thomas Linner (2016), Kamel Saidi and coauthors (2016), and Borja Garcia de Soto and Miroslaw Skibniewski (2020) are merely focusing on the technical evolution since the 1980s predominantly in Japan, and to a lesser degree in the United States.

51 Example works reviewed in Section 1.2.

looked at the effects of robotic systems on work in manufacturing, healthcare, and domestic work settings, mainly through ethnographic studies.[52] The ethnographic study of the work at Canvas, which analyzes worker–robot interactions and the technical labor around Canvas's drywall finishing robot, is an early case-based analysis that expands these existing studies.

Finally, the dissertation makes methodological contributions. The custom-built research tools and their integration with the ethnographic research in the dissertation demonstrate one way in which design and computing may support ethnography.[53] This methodological approach, which is distinct and rarely employed in computational design, is a departure from the field's prevalent focus on computation for building design toolmaking, towards computation for investigating the work of toolmakers and building design professionals. For HCI and CSCW, this approach suggests an alternative to the conventional design–ethnography relationship in these fields, in which the method is typically used to derive technology design implications. The alternative is to use design and computing to serve ethnography.[54] The ethnographic studies framed by this "design (and computing) in ethnography" approach and the research tools in this dissertation add to precedents produced at the intersection of computational design, HCI, and CSCW. Because ethnography is a time- and cost-intensive method, collective knowledge on this method grows more slowly. This makes each extensive fieldwork in any field a precious learning source for researchers. That is the case of my fieldwork, which took more than three years from planning to becoming chapters in this dissertation and will inform future ethnographies of technology design in robotics and construction.

1.6 Limitations

The historical and ethnographic methods I employed in the dissertation come with inherent limitations. The findings of historical research go as far as the researcher's access to documents and people. The documents and people that remained out of reach are this investigation's main limitation and, perhaps, the scope of future investigations. Access to sites, artifacts, and people are equally important in ethnographic research. In the United States, only a limited number of labs and companies focusing on construction robotics, and even a smaller number of commercial-grade construction robots, are available to conduct the investigation I did. Not having access to all these sites or technologies narrows down the options further. As part of this investigation, I engaged with two prominent construction robotics practices, which provided ample findings on which to build a dissertation. That being said, the questions in this dissertation may certainly be extended to other sites in the future.

---

52 Example works reviewed in Section 1.2.

53 "Design (and computing) in ethnography" is a conceptual term I use to describe a particular approach among a group of previous ethnographic work, in which the researcher(s) expand ethnography's capacities on site by in-situ prototyping or discursive design strategies. A notable example is Daniel Cardoso Llach's (2015, 2019) ethnographic study of software and construction in Abu Dhabi, in which he built and used a customized software to record and visualize the 3D modeling conflicts across teams. Another example is Mika Yasuoka's (2009, 2015) study of interdisciplinary collaboration in software development, where she incorporated a conversation analysis tool to pre-process conversational data and visualize the relations between participants' key expressions. A different example is Vera Khovanskaya and collaborators' (2017) ethnography of technology use in families, where the researchers use speculative design conversations as an ethnographic research instrument. These conversations, in which designer and ethnographer perspectives align or contrast around imaginary products, helped researchers filter ideas and gain further understanding of participants' daily practices.

54 This notion of design and computing to serve ethnography next to ethnography informing design and computing is based on the broader discussion initiated by anthropologists Keith Murphy and George Marcus (2013) on making design-ethnography relationship mutual. The anthropologists' work focuses on how design research methods may support the practice of ethnography—next to ethnography informing the practice of design (Murphy & Marcus, 2013). The aim in this dissertation is to bring the anthropologists' notion into ethnography of design and computing. Different than the anthropologists' work, and similar to Cardoso Llach's (2015, 2019) and Yasuoka's (2009, 2015) work, the focus is to leverage (technology) design and computing—rather than design research methods—in ethnography, and develop customized research tools.

Lastly, any case-based research that involves naturalistic observations of sites and participants is scoped through a negotiation between coverage, depth, and time. Although both historical and contemporary sites of construction robotics in the United States are limited, I was fortunate to be on the CMU campus, which hosted two of them, and to be given access to a third in the industry. For this dissertation, I found a balance between coverage, depth, and time by limiting the investigation to three cases over a period of three years from the day I began to work on arranging the sites to the day I finalized the writing. I worked with more than a hundred participants, including construction workers, academic researchers, students, and engineers. However, if given more time, the coverage of this research may have been extended by exploring the questions on a longer list of sites. Alternatively, the depth of investigation may have been furthered through longer fieldwork. Such longitudinal and deeper reach will hopefully be achieved in future work by myself or other researchers building on this dissertation.

1.7 General Outline

In the following chapter, I document my research on the two intertwined historical efforts in robotics at the CMU-RI and the DCE in the 1980s. I first discuss how robotic arms, computer control, and early high-level robot programming languages raised hopes for upgrading the US manufacturing industry, and how these hopes, alongside the concerns on the US–Japan economic race, gave rise to a productivity discourse at the CMU-RI motivating the development of manufacturing systems. I then trace the emergence of construction robotics at the DCE, and show that the field was built not only on technical research but also the engineers' aspirations to reconfigure construction work. I show that these aspirations were built through two discursive elements: a productivity discourse for the US construction industry (resembling the one at the CMU-RI) and high-level descriptions highlighting the distinctive aspects of construction and manufacturing domains to underline the challenges towards robotization along with resemblant aspects to make a case for the undertaking.

Through a close look into these discursive elements, systems, along with a taxonomy of construction work and a modular robot design framework produced at the DCE, I discuss the following points. First, the engineers aspired for a transition from hand tools to intelligent machines on construction sites, which was more ambitious than the ongoing integration of those machines to operate the existing machinery in manufacturing cells. Second, the engineers recognized construction work as a complex, layered constitution of site conditions and activities, which they viewed as challenging for robotics. Meantime, some of them posited part of on-site construction work as more mechanical and analogous to manufacturing work than it really is. Third, while their research was mainly purposed for advancing autonomy, sensing, or motion needed to operate machines on sites, it mostly missed out on the contingency of a set of building tasks and workflows on materials, which is critical to inform and evaluate technologies. Lastly, as I navigate the DCE, I reveal that the work at the Department displayed two divergent technology design approaches, one that concerned building future machines and the other that attended to the present integration of construction work and machines, which co-existed but did not converge in the early days of construction robotics.

In Chapters 3 and 4, I focus on the ReAC project, which incorporated an ethnographic study of construction work (that I conducted) into robotic system development work. In Chapter 3, I present the elements of this study, including the empirically rich descriptions of building tasks, the knowledge of zones and interactions, and the workers' perspectives. In Chapter 4, I document my second ethnographic study focusing on the ReAC's system development efforts. I reveal the ways the ethnography of construction drew robotics and construction closer by describing the roles that the elements of the study played: first, bringing immersion, imagination, and conceptual guides next to

technical knowledge into conceiving a system, second, guiding our analysis of work to focus on the specific qualities of activities and ways they are enacted, which allowed us to situate a robot with an alternative support role (tool-hardware delivery) in a specific carpentry task, and third, informing the development of this system.

In parallel, I trace our group's conception of construction work and automation as they were partly informed by the ethnography of construction. I show that members' conception of construction work is built on an understanding of how work is enacted on site and a sense of how workers view robots, which enabled to see tasks as richly layered and recognize parts that workers invest time, use skills, or attach value. I discuss that this conception shaped the way we conceived the system, first by leading us to shift our attention to relatively hidden, ancillary activities on sites suitable for robots, and second, by leading us to separate the menial support activities that might be well-suited to robots, and primary building activities that should be kept in the hands of workers. I situate these moves as the group's alternative approach to the historically prevalent automation ambitions, which is robotically supporting labor while preserving time-invested, skill-invested, and valued parts of work; and their reimagination of the robots' role as instruments of task support in construction.

In Chapters 5 and 6, I present my ethnographic study at Canvas, which examines the development and deployment work around the company's drywall finishing robot Basil. By reviewing the early efforts that originated this machine in Chapter 5, I trace the engineers' conception of drywall finishing work as a layered practice combining steps tailored for human capabilities and steps for those of machines, and their conception of the work as a material-oriented practice contingent on an interplay between action, tool, and material. I discuss that these conceptions guided the engineers to take key departures from some of the prevalent approaches in robotics: from chasing full automation towards searching for an appropriate step for a machine, from emulating worker skills towards focusing on understanding the finishing material and quality, and from overfocusing on advanced levels of autonomy, sensing, or motion towards experimenting with the material process to make the machine's outcome reliable.

By analyzing Basil during a performance test in the same chapter, I show that the machine reshapes part of drywall finishing into an iterative worker–machine sequence mirroring the system's sense-think-act architecture, and the conventional finishing workflow into work-sharing between finishers and the machine. Within these efforts, I identify the emergence of the finishers' knowledge of steps, finishing quality, and quality evaluation in informing Basil's initial development and evaluating its current performance, displaying a notable convergence of robotics and construction knowledge. Lastly, in Chapter 6, I discuss how Basil reconfigures drywall finishing work. Based on an observation of a job site workflow in person and through the workflow visualization tool tracing worker-robot interactions, I reveal the first way is by rearranging the skill set (involving a degree of deskilling and a form of upskilling). The second way, captured by observations of everyday work and interactions around Basil mapped by the network graph visualization tool, is by establishing a new labor that involves distinct interdisciplinary interactions, roles, and expertise attached to the trade.

In the final chapter, I discuss the key findings regarding the linkages and shifts between past and contemporary conceptions of construction work and systems along with the ones between past aspirations to reconfigure construction work, and its current reconfiguration through state-of-the-art systems between the three cases. I present the main arguments along with reflections on technology design and pedagogy stemmed from the dissertation.

Chapter 2

Decomposing Construction Work and Reimagining It With Robots: CMU (1979–1989)

2.1 Outline

The 1980s saw key advancements in robotic arms and mobile robots, which later became the archetypal technologies for on-site construction robots, together with the institutionalization of robotics and construction robotics research in the United States. CMU Robotics Institute (CMU-RI) and the Department of Civil Engineering (DCE) were two prominent actors in this context. In the early 1980s, the Robotics Institute was chasing programmable automation with arm-based manipulators, early high-level robot programming languages, and generous funding from an engineering industry giant called the Westinghouse Electric Corporation. Meanwhile, the Department was building a new field combining construction and robotics, triggered by knowledge transactions with the Robotics Institute, the excitement around the development of two mobile robots by the Department's Civil Engineering and Construction Robotics Laboratory (CECRL), and the ambitions focusing on the robotization of on-site building construction in the United States.

This chapter takes a close look at these two intertwined efforts, based on primary sources including research reports, articles, newspapers, digital collections, and oral accounts along with secondary historical sources. I first briefly outline robotics research in the 1960s and 1970s before the CMU-RI, focusing on the key advancements in robotic arms, such as computer control and the early high-level robot programming languages that unlocked these technologies' real potential for manufacturing automation. I show how these advancements brought about hopes for the US manufacturing industry, and together with the concerns related to the US–Japan economic race, motivated the launch of the CMU-RI. By examining a report by the Institute positioning robots in the US manufacturing industry and a system developed for Westinghouse Corporation, I trace how these hopes and concerns generated a discourse promoting productivity (via computer-controlled arms and high-level languages) and motivating the systems at the Institute.

Next, I turn to the DCE and show that building construction robotics involved not only technical research but also aspirations to reconfigure construction work based on two discursive elements: a productivity discourse (similar to the one at the Robotics Institute), and high-level descriptions highlighting the distinctive and resemblant aspects of construction domain compared to manufacturing. Based on these elements and systems developed at the DCE, I first discuss that the civil engineers situated the same technology, which had shown promise in manufacturing's transition from hard to programmable automation, and a similar discourse emphasizing productivity through the mobilization of this technology, within a work domain, which required a leap from hand tools to intelligent machines in the 1980s. Second, I discuss that the distinctions portrayed by the engineers show their recognition of construction as a complex, layered makeup of site conditions, and similarities drawn by some of them suggests they conceptualized part of on-site construction work as more mechanical and analogous to manufacturing work than it really is. Lastly, I contend that the engineers mostly missed out on the material-oriented character of a set of building tasks and workflows (critical to account for in developing technologies), while pushing to advance autonomy, sensing, or motion needed to perform those tasks with machines.

Among the work at the DCE, I closely analyze two studies aimed at closing the gap between building tasks, sites, and systems, along with the two unconventional technology design instruments that these studies produced: Abraham Warszawski and Dwight Sangrey's (1985) taxonomy of construction work to identify robot performance requirements, and Miroslaw Skibniewski's (1986) modular robot design framework. While the high-level descriptions demonstrate the recognition of construction as a complex, layered makeup of site conditions at the DCE, I identify that Warszawski, Sangrey, and Skibniewski's approach of decomposing work and instruments indicate a conception of

construction work as a complex, layered makeup of activities. Drawing from an analysis of these instruments, I discuss that these engineers focused on the general character and commonalities of building tasks to match each or a set of them with an archetypal robot configuration as a cost-efficient way of robot integration. However, because these instruments also stemmed from a high-level view of construction work like the descriptions, I point out that they overlooked the situated nature and idiosyncrasies of tasks.

As I describe the efforts at the DCE, I identify that construction robotics was situated in two dominant but divergent approaches in its early years. The first one embedded in the CECRL's mobile robot prototypes focused on building technologies that fulfilled the missing capabilities of machines to handle construction work. The second technology design approach at the Department, which is visible in Warszawski, Sangrey, and Skibniewski's studies, concerned identifying, decomposing, and classifying construction tasks for robot development and adoption. While the first one draws a path through robotization through customized chassis, sensors, and models, the second one draws a path with an attempt to pull the knowledge of the practice and of technologies closer. Nonetheless, these two paths ran parallel at the DCE, and the worlds of robotics and construction were seen distant in the sense that the technologies would inevitably disrupt labor.

2.2 The Road to the 1980s

2.2.1 Robotic Arms and Computer Control

The 1960s was a decade of both fear and excitement about technology in the United States. On the one hand, there were the changes in labor caused by ongoing automation in the industries since the 1950s. On the other hand, there was the rise of a novel research field called AI in the nation's prominent institutions, supported by a new governmental entity, the Advanced Research Projects Agency (ARPA).[1] Changes in labor consisted of an increase in unemployment, with a series of recessions in 1953, 1958, and 1960 (Harris, 2020; Mitchell, 2017).[2] After 1958, unemployment reached an alarming seven point five percent, and after 1961, it stood at seven point one percent as a result of large layoffs in manufacturing (Wartzman, 2017; Mitchell, 2017). According to big-name magazines, the culprits for the layoffs were industrial machines. "We are stumbling blindly into the automation era with no concept or plan to reconcile the need of workers for income and the need of business for cost-cutting and worker-displacing innovations," was the criticism voiced by *The Nation* in a 1958 issue (as cited in Wartzman, 2017). "The rise in unemployment has raised some new alarms around an old scare word: automation," claimed "Time" magazine, reflecting on the 1960–1961 recession ("Business," 1961).

This period of recessions and media speculation fueled the government's concerns regarding technology and labor. In 1964, President Lyndon Johnson formed the National Commission on Technology, Automation, and Economic Progress, comprised of economists, engineers, businessmen, and labor union leaders. The Commission was tasked with investigating the speculated connection between unemployment and technological change and, ultimately, helping construct new economic policies (National Commission, 1966). From another perspective, the Commission would hopefully change the negative impression of technology spreading across the

1 For more on DARPA (previously ARPA), see the book by the National Research Council (1999), titled *Funding a Revolution: Government Support for Computing Research*.

2 See more on the recessions in the United States in "US Recessions Throughout History: Causes and Effects" on Investopedia website (2024).

country for a more balanced and proactive view. The mission started with the remarks of President Johnson (1964) as he signed the bill creating the Commission:

> *Technology is creating both new opportunities and new obligations for us—opportunity for greater productivity and progress—obligation to be sure that no workingman, no family must pay an unjust price for progress. Automation is not our enemy. Our enemies are ignorance, indifference, and inertia. Automation can be the ally of our prosperity if we will just look ahead, if we will understand what is to come, and if we will set our course wisely after proper planning for the future.*

The two-year investigation resulted in a report titled "Technology and the American Economy." Featuring statements we have since grown used to hearing, the Commission (1966) in the report's introduction described the public's concern about technology and jobs in those years:

> *Still another concern, perhaps the one most responsible for the establishment of the Commission, has arisen from the belief that technological change is a major source of unemployment. This concern has been fostered by the substantial and persistent unemployment during the period 1954-65. The fear has even been expressed by some that technological change would in the near future not only cause increasing unemployment, but that eventually it would eliminate all but a few jobs, with the major portion of what we now call work being performed automatically by machine.*

In the same section, the Commission (1966) projected that "machines and automated processes" would free humans from "routine and mechanical work" and, in parallel with the President's remarks, redirected attention to "[how] to marshal the needed technologies" as the real challenge. The further sections offered their assessment that exonerated the technology in the case of unemployment. According to this assessment, the unemployment rates since the 1953 recession were not a consequence of fast-paced technological change but a function of four explicable variables: productivity increase, labor force growth, average annual work hours, and the total production of goods and services—also known in economics as the gross national product (GNP)" (National Commission, 1966). In six of the preceding twelve years in the United States, the GNP did not grow faster than the labor force and productivity (the output of goods and services per man-hour), while the average annual work hours did not change dramatically (National Commission, 1966). A faster-growing GNP, if achieved, would potentially create new jobs, which would compensate for unemployment due to increased productivity and employing the new members of the labor force. Thus, the Commission (1966) concluded that the high unemployment rates were not an "inevitable consequence" and were even avoidable should the variables be more effectively controlled via essential policy measures.

While the government was busy healing unemployment, a new strand of research, which would become an increasingly important part of the technology-labor discussion, was in the making. Started as a scientific inquiry among a few select scientists in the 1950s**,** AI research was taking a more organized shape in the 1960s at the nation's prominent institutions (Buchanan, 2005; National Research Council, 1999).[3] At Carnegie Tech, later re-established as CMU in 1967, a research

3 A couple key events in the 1950s described in "Funding a Revolution" by the National Research Council (1999) are "a conference on automata studies" in 1952 organized by Claude Shannon, who was then working at Bell Laboratories, and John McCarthy, who was then at the Mathematics Department at Princeton; as well as the renowned "The Dartmouth Summer Research Project on Artificial Intelligence" in 1956, initiated again by Shannon and McCarthy, who was then at the Mathematics Department at Dartmouth College, along with Marvin Minsky, who was then a fellow at Harvard University, and Nathaniel Rochester, a researcher at IBM.

stream was being built around Herbert Simon and Alan Newell's work at the Graduate School of Industrial Administration since the mid-1950s (National Research Council, 1999). At MIT, a similar coordinated effort was the Artificial Intelligence Project founded by John McCarthy and Marvin Minsky in 1957, which went on to become the MIT Artificial Intelligence Laboratory (MIT AI Lab) (National Research Council, 1999).[4] After McCarthy moved to Stanford and launched the Stanford Artificial Intelligence Laboratory (SAIL) in 1963, the lab joined the other two as the center of AI research on the West Coast (National Research Council, 1999). Last but not least was the Artificial Intelligence Center, which was founded in 1966 under the Stanford Research Institute (SRI) (National Research Council, 1999). These four centers constituted the major intellectual force in AI research in the 1960s (Fleck, 1982).

The other and perhaps more significant element that escalated AI research in the 1960s was the Department of Defense's support. In 1958, the Department established ARPA with the aim of facilitating scientific research towards military technologies (DARPA, 2023). In the 1960s, ARPA selected AI as a concentration for the advancement of military systems, and launched its Information Processing Techniques Office, which began to heavily fund AI research at Carnegie Tech, MIT, Stanford, and SRI in 1962 (National Research Council, 1999). ARPA dispensed seventy-five percent of the AI research funding in the United States between 1964 and 1974, and a chief portion of that money went to the four AI centers at these institutions (Fleck, 1982). Conjoining intellectual and financial forces, these centers would lead in shaping AI research and technologies in the 1960s and beyond. In contrast with the fears of the decade, the pioneering work coming out of these centers carried optimism and excitement about technology and change.

In this context of technology and labor in the 1960s, we see the first seeds of robotics research in universities, planted chiefly at SAIL and MIT AI Lab. From the 1960s and to the late 1970s, two key advancements, computer control and high-level robot programming languages, that I overview in this chapter turned arm-based manipulators that were clunky, task-specific, and time-consuming to program into a technology of hope for the US manufacturing industry. By promising a transition from hard to programmable automation, computer-controlled arms together with high-level languages raised hopes for the future competitiveness of the industry, and became the subject of a discourse on productivity prevailing in the country and the CMU-RI as I describe in this chapter.

The general landscape of robotics in the 1960s, as Bernard Roth (2016)—a then young assistant professor at SAIL—recalls, involved SAIL and MIT AI Lab being "the only sizeable academic ventures" among "very little academic or commercial activity."[5] Robots at the time were mainly associated with industrial purposes; in fact, the first reliable ones in the United States were arm-based manipulator systems for simple manufacturing tasks. The renowned example was Unimate, developed by the engineer-entrepreneur duo George Devol and Joseph Engelberger and later deployed on a General Motors (GM) assembly line for pick-and-place and welding in 1961 (Figure 1).[6] Real-world applications as such came first as precursors to robotics research at SAIL and MIT AI Lab, and in the field's formative years, arm-based manipulation was the primary concentration of researchers in these centers. Distant from real-world manufacturing settings, the researchers in

4 Two different sources state different dates for the launch of the Artificial Intelligence Project at MIT. While National Research Council's (1999) book *Funding a Revolution* states 1957, Encyclopaedia Britannica's "Marvin Minsky" page written by Michael Aaron Dennis (2024) states 1959.

5 See Bernard Roth's (2016) foreword in *Springer Handbook of Robotics* edited by robotics professors Bruno Siciliano and Oussama Khatib (2016).

6 See more information about Unimate in the Henry Ford Museum (2022) digital collections available online or the article by Allison Marsh (2022) titled "In 1961, the First Robot Arm Punched In," published online in IEEE Spectrum magazine. For more technical information on Unimate, see Joseph Engelberger's (1980) book *Robotics in Practice: Management and Applications of Industrial Robots*.

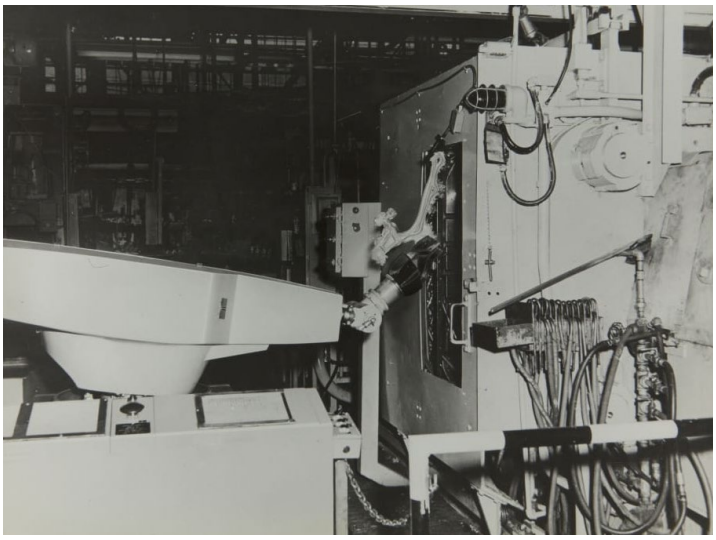


Figure 1. Unimate at the General Motors factory (1961-1962) (Source: The Henry Ford)

these centers acquired or built arms as vessels to achieve a significant research ambition: making these machines more easily and efficiently programmable with a computer. The key developments in this endeavor came from SAIL as I review in this section.

By the early 1960s, Unimate, as an electronically controlled basic programmable arm, was essentially at the frontier of these technologies, which quickly shifted over the next few years. To program Unimate's motions, engineers would store pre-programmed sequences on its rotating drum memory and run these stored sequences to execute its motions. The arm was hydraulically actuated without any position sensing capacity, which made it difficult to maintain full accurate control. Two years after the installation of Unimate at the GM factory, SAIL researchers acquired the Rancho Arm, from the Rancho Los Amigos Hospital, where the arm was originally built as an assistive technology for handicapped patients.[7] The researchers interfaced this electrically actuated arm with a Programmed Data Processor Model 6 (PDP-6) mainframe computer and started upgrading it meticulously to push the state of the art in computer control and sensing. In a relatively short timeframe, this arm, wittingly renamed "Butterfinger", surpassed Unimate with improved programmability and sensing features.

A pick-and-place demonstration by researchers Karl Pingle, Jeff Singer, and Bill Wichman in 1967 showcased these features.[8] For this demonstration, the researchers integrated a camera into the PDP-6 mainframe for capturing toy-size cubes on a black backdrop and deployed a vision and arm-control program they wrote on Butterfinger. Through this system, they demonstrated rudimentary image-based sensing and programmable motion control in a pick-and-place workflow, in which the program detected the blocks and the arm picked and stacked them by size (Figure 2).[9] Aside from image-based sensing, the other notable improvement was Butterfinger's precision in positioning and grip during the workflow. This was enabled by the potentiometers of the arm, providing the researchers feedback to know the position of joints in the workspace. The work around Butterfinger

7 The information about the Rancho Arm in this section is based on a video demo produced by Gary Feldman (1967), accessible online on MIT CSAIL Video Archive. See the source in Bibliography.

8 The video demo by Gary Feldman (1967) features this pick-and-place experiment.

9 According to the description in the demo video, the system included a camera sending images of the scene to PDP-6, which would run the visual processing program named the "Edgefinder" on the images. The Edgefinder would find the corners of each block, detect the block by size, compute its position, and the arm control program would execute the stacking sequence.

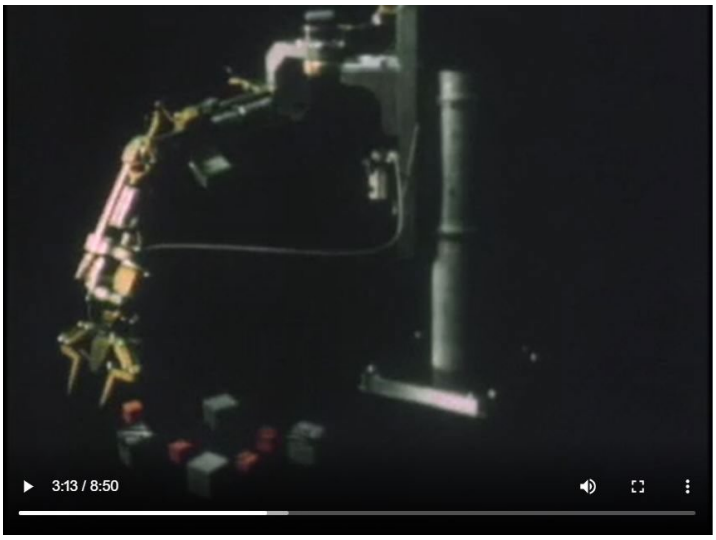


Figure 2. Butterfinger during the pick-and-place demo in 1967 (Source: MIT CSAIL)

was only the beginning of computer control and sensing experiments in universities. By the mid-to-late 1960s, a series of arms, such as the Orm and the Hydraulic Arm at SAIL as well as the Minsky Tentacle Arm at MIT AI Lab, were built and used in similar experiments.

Yet the most significant milestone in computer control and sensing was reached in the late 1970s with the Stanford Arm. This "first truly functional arm" at SAIL, as Bernard Roth (2016) puts it, was primarily developed by a SAIL Ph.D. student, Victor Scheinman, with a control software written by another SAIL Ph.D. student, Richard Paul. The arm was an important artifact in the progression of arm-based manipulation as it facilitated and marked the beginnings of time-efficient computer control. It can be considered Scheinman's response to the difficulties he experienced firsthand in programming both Butterfinger and the Hydraulic Arm with the lab's time-sharing system (Figure 3). In the 1960s, Scheinman and his peers had to program these robots over and over for experiments, and the joint control of this hardware required excessive time and computing in each round. Years later, in an email to the Computer History Museum, Scheinman (2011) still recalled programming the Hydraulic Arm by running the PDP-6 with a "boost" feature that SAIL researchers installed to run the 1960s game Spacewar!.[10] Motivated to cut these long tedious hours of programming, Scheinman along with Bernard Roth and Jahangir Rastegar (1974) started from kinematics and tackled a key question "not previously treated."[11]

The question that Scheinman, Roth, and Rastegar (1974) regarded as "not previously treated" and that became instrumental in shaping the Stanford Arm's kinematic configuration was presented in 1973 at the first Symposium on Theory and Practice of Robots and Manipulators: "for any given manipulator, how many configurations yield the same position and orientation of the terminal device?" This was an important question because its answer, as the researchers (1974) mentioned, would make it possible to find "all possible [arm] configurations" and "determine which configuration

10 The "Spacewar Mode" that Victor Scheinman (2011) mentioned in his letter has also been mentioned in "The Autobiography of SAIL" a short reflection about the Stanford Artificial Intelligence Lab by Stanford University InfoLab (1991) available online.

11 "A question of both theoretical and practical interest that does not seem to have been previously treated" were the exact wording in the 1973 lecture that presented this question. The lecture is titled "On the Design of Computer Controlled Manipulators," it was presented in the "Symposium on Theory and Practice of Robots and Manipulators," and later published as Roth, Rastegar, and Scheinman (1974) within the symposium proceedings.

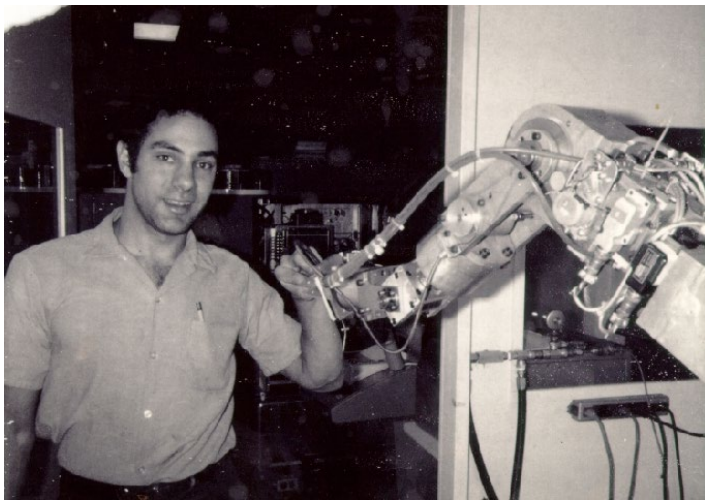


Figure 3. Scheinman next to the Stanford Hydraulic Arm (Source: Stanford University InfoLab)

is best" to execute a motion state. This would enable, in mathematical terms, a "closed-form solution," that is, the calculations of joint position and orientation could be solved in a finite number of operations more simply and faster (Roth et al., 1974). Compared to an iterative solution, this meant time-efficient computing of arm motion sequences, and the researchers (1974) consciously determined the Stanford Arm's kinematic configuration to attain this.

The Stanford Arm added more than just time-efficient computing to its predecessors. The arm was electrically actuated with direct current motors installed in its six joints, which eliminated the issues of jolts and accuracy in speed control. It was also equipped with additional sensors providing joint velocity, joint position, and force/torque feedback to its control software. The control software was architected into an "arm-control servo" and a "trajectory-planning" component, which allowed to dedicate a real-time PDP-6 and a time-shared PDP-10 computer respectively to each of them (Earnest et al., 1973). Trajectory planning was written in a more intuitive high-level language (Earnest et al., 1973). To plan a trajectory, the first step would be to perform an early form of learning from demonstration (LfD) by moving the arm's gripper, also called "the hand," to its various end positions (Bolles & Paul, 1973). After storing these positions, the next step would be to run the program to generate trajectories, and based on these trajectories as well as the velocity and position feedback, the arm-control servo would set the joint output torques needed to execute the motion (Bolles & Paul, 1973).

The advances brought by the Stanford Arm in kinematics, sensing, and control were fundamental to the development of what the present-day robotics community categorizes as a new generation of industrial robots throughout the 1970s (Zamalloa et al., 2017).[12] Specifically, these advances were the closed-form arm solution for computational efficiency, the use of electric actuators and servo-controllers over hydraulic actuators (as seen in the shift from the Hydraulic Arm to the Stanford Arm), and basic sensing capabilities in vision, velocity, position, and force/torque crucial to motion planning and control. The driving force behind these progressions was to make robot programming easier and more efficient with computers. Examples of this new generation of robots that benefited from these progressions included the Vicarm developed by Scheinman in 1973, later bought by Unimation, which became the foundation for a commercial-grade arm called Programmable Universal Machine

12 This technical categorization created by Irati Zamalloa and colleagues (2017) was used also by Alessandro Gasparetto and Lorenzo Scalera (2019) in their historical overview of industrial robots.

for Assembly (PUMA), developed in collaboration with GM in 1978 (Gasparetto & Scalera, 2019; Marsh, 2022). Another significant example was the Tomorrow Tool (T3) developed by Cincinnati Milacron in 1974 as a commercially available computer-controlled robot (Gasparetto & Scalera, 2019; Wallén, 2008).

The programming features of these new generation machines made them conveniently reprogrammable and reusable for different tasks, and therefore, appealing to the manufacturing industry. PUMA was deployed on GM car part assembly lines in the United States, whereas T3 was used in Ford and Volvo facilities in Australia and Sweden, respectively (Marsh, 2004; The Powerhouse Museum, 2019; Gasparetto & Scalera, 2019).[13] At the end of the 1970s, a PUMA would be programmed by "a primitive motion level" language called Variable Assembly Language (VAL), and a T3 with a "point-to-point level" language named after the arm (Shin & Bonner, 1982; Srihari & Deisenroth, 1988). In VAL, one could write loops, branching statements, and subroutines with close-to-common-language commands, debug quickly, and keep a program as an executable code that could be run directly on the computer (Shin & Bonner, 1982; Srihari & Deisenroth, 1988). In T3, one could teach the arm a task sequence by demonstration, modify the program commands for any corrections, and keep the programmed sequence for reuse (Shin & Bonner, 1982). All these were important contributions to ease the reprogramming of these machines for various tasks.

From the 1960s to the late 1970s, arms in the United States moved from hydraulic actuators to electric actuators and closed-form solutions, accompanied by rudimentary sensor feedback for improved motion planning and control. More importantly, the way these robots were programmed shifted from pre-programmed sequences on rotating drum memory to computer control, as in the research experiments at SAIL. In the late 1970s, this became less laborious with high-level languages, as in the new generation of arms deployed in the auto industry. These were significant advancements because, for nearly two decades, arms were essentially task-specific devices that were excruciatingly time-consuming to program. The real potential of these technologies in the imagination of industrialists and researchers was that they could efficiently switch between various manufacturing tasks, and that potential was yet to be unfolded. At the turn of the 1980s, this began to change, and computer control and high-level languages created a new wave of excitement. The field of robotics and the CMU-RI would surf this wave in the 1980s.

2.2.2 The Launch of the Robotics Institute

The establishment of the CMU-RI in 1979 was partially the outcome of Richard Cyert's managerial mindset in administrating CMU as its president. It was also partially the outcome of the pressure put on Cyert at a meeting in Washington DC, where the government voiced concerns about Japan's rise in the global economic race through industrial robot adoption. Cyert (1978) saw research as "part of the product package [of private universities]" to be leveraged for financial "equilibrium." A dedicated robotics institute could possibly be contributory to CMU in this way that he envisioned. Such an institute could also possibly aid US manufacturing by developing new technologies through industry partnerships, and therefore, contribute to the nation's economy in the global race against Japan. In a 2018 interview, Raj Reddy, who would become the Robotics Institute's founding director, recalled Cyert's return from Washington DC and checking in with the computer science faculty about the existing research in robotics (Markoff, 2018). As Reddy quoted Cyert in the interview, "Hey, I came

13 Marsh (2004) also notes that the arm "failed to revolutionize lightweight assembly work." The exact wording of the Powerhouse Museum (2019) is "thought to have been used in a training facility" at the plant and "from around 1980 until the plant's closure in 1994." The fact that the arm was used in Volvo plants in Sweden is based on Gasparetto and Scalera's (2019) historical overview.

from Washington, DC. There's a lot of concern that the Japanese are eating our lunch with robotics, this, that, and the other. Do we have anything to show for it now in that area?" Cyert asked him and Joseph Traub, the head of the Computer Science Department, and Alan Newell, a faculty member well-renowned in AI circles (Markoff, 2018).

Reddy, Traub, and Newell's answer to Cyert's inquiry was basic research in computer vision and speech recognition based on a million dollars a year (Markoff, 2018). For instance, Reddy, who started at SAIL with a Ph.D., had been involved in AI and robotics research for more than a decade by then. SAIL's first computer, PDP-1, allowing speech input got Reddy interested in researching speech recognition for his Ph.D., which he finished in 1966, and building a speech-based interface for the Stanford Arm later as an assistant professor before joining CMU in 1969 (Hsu, 2021).[14] His projects in the 1970s were a continuation of this work, and according to Reddy, the "closest thing [to robotics]" at CMU (Hsu, 2021).[15] The three professors advised Cyert that a ten or a hundredfold increase in funds would be needed to "have [a] real impact on the world" (Markoff, 2018). As an economist and a former member of the Navy, Cyert knew where to look for this money. With Reddy, he began to knock on the doors of the manufacturing industry and the Army.

Cyert and Reddy's first destination was Westinghouse Electric Corporation, an almost century-old industry giant at the time with headquarters located a short drive away from CMU (Markoff, 2018). Having established a list of manufacturing plants in multiple states and abroad, the company specialized in the manufacturing of electrical equipment and large-scale production infrastructure.[16] It was Tom Murrin, the president of the company's Energy and Advanced Technology Group, whom Cyert invited to support CMU's ambitious undertaking. Murrin was as a follower of Japanese manufacturing according to various sources, and his and Cyert's minds were aligned on what Japan did better than the United States in industrial production.[17] As Cyert (1981) stated in a Joint Economic Committee Congress Hearing, Japan was "not ahead of [the United States] either technologically or scientifically," but the gap concerned the lack of funding and of a streamlined research effort in the United States.[18] Both presidents believed in university–industry partnerships to close that gap and re-establish the country's competitive edge as Reddy (1989) mentioned in retrospect at the Institute's tenth year anniversary. Based on this premise, Westinghouse agreed to kick-start the Robotics Institute with five million dollars over the following five years (CMU Robotics Institute, 2022).

The other destination Cyert and Reddy went to in search of funds was the Pentagon (CMU Robotics Institute, 2022). Cyert, who served during World War II, was no stranger to the Department of the Navy and its research arm, the Office of Naval Research (ONR), founded in 1946 (Ravo, 1998; Babb, 2022). Active since, this military agency had been funding university and industry research in fields such as biology, physics, and robotics as well as defense technologies and was a perfect

14 For more on the computers at SAIL, see the online description titled "SAIL Away" by Lester Earnest (1995), who was a member of the lab in the 1960s and 1970s.

15 Reddy joined the Computer Science Department at Carnegie Mellon University in 1969 (Hsu, 2021).

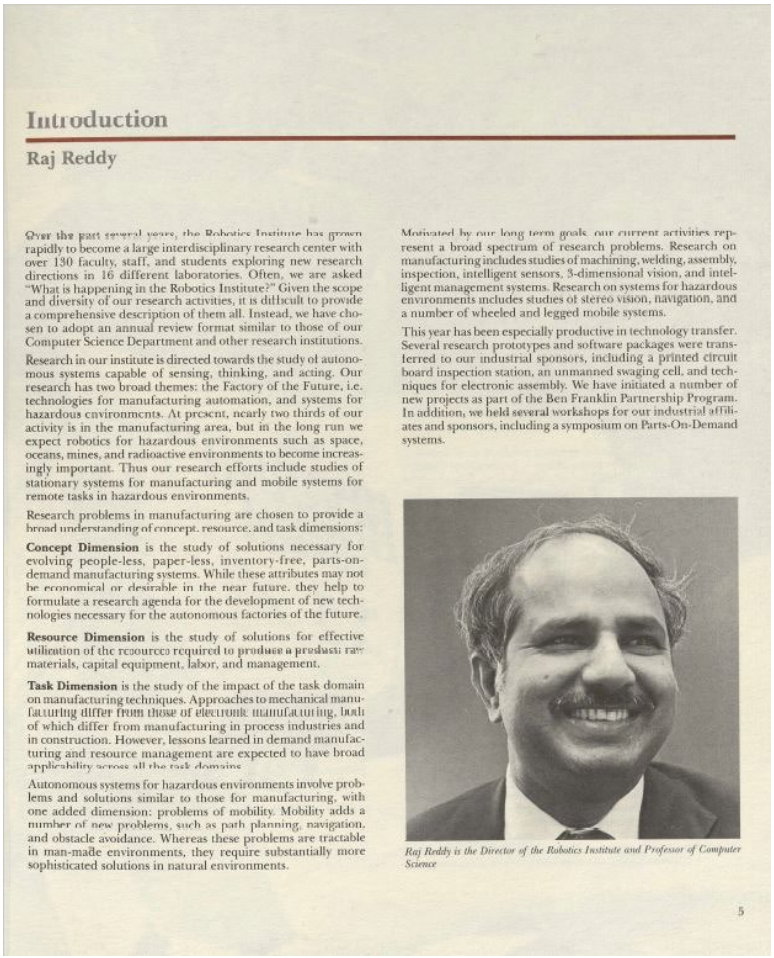
16 For more information on Westinghouse Electric Corporation, see the article titled "The Westinghouse Electric and Manufacturing Company" published on the Library of Congress website (2023), and "Westinghouse Electric Corporation" page in Encyclopaedia Britannica (2023).

17 See for example, Raj Reddy's (1989) introduction to the *Robotics Institute 1989 Annual Research Review* titled "Looking to the Future," and the short article by Jason Togyer (2012) titled "RI Co-founder Tom Murrin: 1929-2012," published on Carnegie Mellon University School of Computer Science website.

18 The title of this 1981 Hearing (Congress of the United States, 1982), where Cyert (1981) presented a statement, is "Japanese Productivity: Lessons for America." In his statement, Cyert (1981) elaborated on why Japan was "not ahead [of the United States]" with an interesting point: "The sad aspect is that the knowledge and, in many cases, the robots themselves are American. They have been adapted in so many ways by the Japanese, but more importantly they have been adopted by Japanese firms."

potential partner for CMU's plans with its interest in autonomous navigation research (Babb, 2022). ONR's support was already so large by the late 1970s that it encompassed a long list of Nobel laureates, including CMU's well-known Herbert Simon.[19] The chief of ONR, Admiral Albert Baciocco, was the agency's main connection with its academic partners in those years. Murrin was an avid AI and robotics proponent in the manufacturing industry, and Baciocco was his equivalent in military circles. Cyert's (1981) statement emphasizing the university–industry partnership at the Joint Economic Committee Congress Hearing also included praise for Baciocco's initiation of AI lectures at the Pentagon. A few meetings between Cyert, Reddy, and Baciocco secured a million dollars of funding for the Robotics Institute to conduct autonomous navigation research (Hsu, 2021).[20]

In October 1979, CMU launched the first dedicated robotics institute in a US university, with Raj Reddy as its first director (Figure 4). The partnership with Westinghouse made manufacturing automation a major line of research on which the Institute would continuously work throughout the 1980s. ONR's additional funding positioned autonomous navigation as the second major line the Institute would advance significantly in the following decade. Separating robotics research from the bureaucratic web of computer science and engineering was an unprecedented move for the



Introduction

Raj Reddy

Over the past several years, the Robotics Institute has grown rapidly to become a large interdisciplinary research center with over 130 faculty, staff, and students exploring new research directions in 16 different laboratories. Often, we are asked "What is happening in the Robotics Institute?" Given the scope and diversity of our research activities, it is difficult to provide a comprehensive description of them all. Instead, we have chosen to adopt an annual review format similar to those of our Computer Science Department and other research institutions.

Research in our institute is directed towards the study of autonomous systems capable of sensing, thinking, and acting. Our research has two broad themes: the Factory of the Future, i.e. technologies for manufacturing automation, and systems for hazardous environments. At present, nearly two thirds of our activity is in the manufacturing area, but in the long run we expect robotics for hazardous environments such as space, oceans, mines, and radioactive environments to become increasingly important. Thus our research efforts include studies of stationary systems for manufacturing and mobile systems for remote tasks in hazardous environments.

Research problems in manufacturing are chosen to provide a broad understanding of concept, resource, and task dimensions:

Concept Dimension is the study of solutions necessary for evolving people-less, paper-less, inventory-free, parts-on-demand manufacturing systems. While these attributes may not be economical or desirable in the near future, they help to formulate a research agenda for the development of new technologies necessary for the autonomous factories of the future.

Resource Dimension is the study of solutions for effective utilization of the resources required to produce a product: raw materials, capital equipment, labor, and management.

Task Dimension is the study of the impact of the task domain on manufacturing techniques. Approaches to mechanical manufacturing differ from those of electronic manufacturing, both of which differ from manufacturing in process industries and in construction. However, lessons learned in demand manufacturing and resource management are expected to have broad applicability across all the task domains.

Autonomous systems for hazardous environments involve problems and solutions similar to those for manufacturing, with one added dimension: problems of mobility. Mobility adds a number of new problems, such as path planning, navigation, and obstacle avoidance. Whereas these problems are tractable in man-made environments, they require substantially more sophisticated solutions in natural environments.

Motivated by our long term goals, our current activities represent a broad spectrum of research problems. Research on manufacturing includes studies of machining, welding, assembly, inspection, intelligent sensors, 3-dimensional vision, and intelligent management systems. Research on systems for hazardous environments includes studies of stereo vision, navigation, and a number of wheeled and legged mobile systems.

This year has been especially productive in technology transfer. Several research prototypes and software packages were transferred to our industrial sponsors, including a printed circuit board inspection station, an unmanned swaging cell, and techniques for electronic assembly. We have initiated a number of new projects as part of the Ben Franklin Partnership Program. In addition, we held several workshops for our industrial affiliates and sponsors, including a symposium on Parts-On-Demand systems.

Raj Reddy is the Director of the Robotics Institute and Professor of Computer Science

5

Figure 4. Raj Reddy in the *Robotics Institute 1983 Annual Research Review* (Source: *The Robotics Institute 1983 Annual Research Review*).

19 A full list of Nobel laureates sponsored by ONR can be found on ONR's website (2023).
20 In this interview with Hansen Hsu (2021), Reddy mentioned this funding as "a million dollars a year."

institutionalization of robotics.[21] Whereas robotics was only a branch at SAIL and MIT AI Lab then, CMU turned it into a standalone entrepreneurial project with dedicated resources. The excitement of the late 1970s around the deployment of PUMA and T3 in the auto industry and the robotic arms' expanding task versatility with reprogramming were coupled with aspirations for manufacturing automation and autonomous navigation, which arose at CMU's brand-new Institute. Backed by two large funds, the Institute was ready to chase these aspirations.

2.3 From Hard Automation to Programmable Automation: Robotics at the Institute

In the next two sections, I delineate the technology-labor context within the 1980s as the first decade in which the Robotics Institute operated. As I will expand further in the sections, this decade saw a decline in manufacturing employment due to high productivity, an increase in robot adoption and high-level languages in manufacturing, and the competition in global trade against Japan, which was a motivation behind the Institute's launch. Computer-controlled arms and high-level languages continued to promise a transition from hard to programmable automation, and raised hopes for salvaging US manufacturing via an increase in productivity, product diversity, and exports. As I noted earlier in this chapter and detail in these sections, they became the subject of a productivity discourse prevailing in the country and the Institute. I show these through analyzing two major works produced at the Institute: a report, envisioning higher batch production output at a lower cost and with less labor, facilitated by robots, and a manufacturing automation system developed for Westinghouse Corporation.

2.3.1 High-Level Languages, Reprogrammability, and the High Hopes for Manufacturing

The first decade of the Robotics Institute started with a tumultuous couple of years for the US manufacturing industry. As a result of the growing number of jobs in the 1960s and 1970s, the manufacturing employment rate was at an all-time high with twenty-two percent of total nonfarm jobs only a few months before the Institute's launch in 1979 (Harris, 2020). This peak employment rate ended quickly when the 1980 and 1981–1982 recessions cost the industry a total of nearly four million jobs (Plunkert, 1990). Despite subsequent partial recoveries, a decline in the share of total employment combined with a productivity increase, which kept the total output stable in spite of labor losses, circumscribed the industry's fate in the 1980s (Plunkert, 1990). A good illustration of this decline was the steel industry, part of which was housed in the hometown of the Institute, Pittsburgh. The industry suffered one of the highest numbers of job losses, and this decrease in the number of jobs and plants almost doubled its productivity rate (Plunkert, 1990).

Grappling with recessions and a decline in its manufacturing industry, the United States was losing pace against its global competitors. The strongest rival threatening its lead in the global economy was Japan, which, in contrast, was enjoying strong economic growth primarily due to its fast-developing manufacturing industry. Technology and productivity levels in Japanese factories in the early 1980s allowed the country to achieve higher outputs with low production costs and low price tags, leading to Japan's strong stance on both domestic and foreign markets (Komiya & Irino, 1992). Producing more and cheaper for foreign markets gave Japan a significant edge over the United States in international trade, but its domestic market being relatively closed created "trade friction" between the two countries (Langdon, 1983). In January 1982, Japan announced an enormous

21 The traces of this move can be seen in Takeo Kanade's (1987) distinction of robotics as a field in his introduction to the *Robotics Institute 1987 Annual Research Review*. Kanade is an early member and later director of the Institute. Recalling 1979 and the launch of the Institute, Kanade in this article states: "We recognized that this challenge [of understanding how to create autonomous machines that could sense, think, and act] would require a discipline closely related but distinct from computer science and artificial intelligence, mechanical, electrical and other engineering fields."

increase in its trade surplus with the United States, which reached more than thirteen billion dollars; in response, the Reagan administration demanded greater access to the Japanese domestic market for US exports (Farnsworth, 1982). This occurred just a few months after Japan's auto industry exports, which dominated the US market, were imposed a three-year cap based on an agreement between the two governments (Tharp, 1981).

Despite the decline in employment, on the technology side, the manufacturing industry was doing well in terms of robot adoption. In parallel with the upward trend in employment, the number of active robots in the industry was also increasing in the 1970s. Surveys conducted in different years by various stakeholders, such as the market research company Frost and Sullivan, manufacturing equipment company Prab Conveyors, and even *The American Machinist* magazine showed an impressive rise from two hundred to three thousand machines (Ayres & Miller, 1981). The high-level languages, which made robots easier to operate, and deployments of the late 1970s fostered demand, and the increase in the number of machines accelerated as we entered the 1980s. The year 1980 saw the beginning of employment decline and an increase in robot production equivalent to more than a third of the entire 1970s (Ayres & Miller, 1981). By the end of 1980, approximately four thousand machines including arm-based manipulators and detection systems had found their place on machining, casting, or inspection lines (Ayres & Miller, 1981).[22] Robots' momentum was so big that the number would quadruple over the following five years (Miller, 1989a).[23]

In parallel with the increasing robot deployments in manufacturing, high-level languages stemming from companies began to allow factories to customize the arms for different parts of assembly lines. Task versatility—these technologies' desired potential for manufacturing—began to unlock. These languages continued to grow in number and capabilities throughout the 1980s. For instance, practical robot reprogramming capability by demonstration or textual languages enabled with VAL, T3, and other manipulator-level commercial languages in the late 1970s expanded with the release of new languages. By 1986, VAL and T3, the most used languages in the industry, were joined by languages such as AML, AR-BASIC, KAREL, and RAIL (Blaha et al., 1986). Developed by IBM, American Robot, GMF Robotics, and Automatix, respectively, these manipulator-level languages had high flexibility so they could be used adaptively for building a range of applications and extensibility as they enabled data abstraction to integrate new robots or sensors (Blaha et al., 1986). These features enabled to use robots more flexibly in manufacturing.

As US manufacturing was slowing down while the production of robots and languages was going up, these technologies became increasingly synonymous with changing the industry's fate and restoring its place in international trade vis-à-vis Japan. Robot production and installation turned into an arms race between the two nations. "Ultimately, however, a nation's success at home and abroad depends upon productivity, upon the quantity and quality of goods turned out per production-line worker. And that is the most important gift of the robot" were the words of Henry Scott Stokes (1982), the chief of the Tokyo bureau of the *New York Times*, summarizing the leverage these technologies could provide in the United States–Japan competition. With university research as well as robot and high-level language production in the industry, the United States was ahead of Japan in technology development. In fact, Japan's robot production relied mainly on licensed technologies from the United States (Stokes, 1982).

22 Note that the United States had a particular definition of robots in those years. See the United States' definition in Richard Miller's (1989a) "Robots in Industry" chapter of his book *Industrial Robot Handbook: Case Histories of Effective Robot Use in 70 Industries*.

23 According to Miller's (1989a) "Robots in Industry" chapter, the Robotic Industries Association (RIA) approximated sixteen thousand industrial robot installations in 1985.

Nonetheless, the country's manufacturing industry was worryingly behind Japan's in terms of robot adoption, and the existing robots were in the hands of only a few auto manufacturing companies (Cyert, 1981; Miller, 1989b). Despite the aforementioned spike in industrial robot installations in the early 1980s, there were nearly three times as many machines operating in Japanese factories at the time (Ayres & Miller, 1981; Stokes, 1982).[24] The total production by robots was estimated to be four times greater in Japan than in the United States at the end of the decade (Stokes, 1982).[25] Other than more reasonable capital costs, the relatively more positive attitude of workers and labor unions towards these automation technologies in Japan was one of the reasons behind this gap (Lohr, 1981). The other pressing issue was that the existing robots were owned and deployed by only a few stakeholders and branches of the manufacturing industry. By the end of 1980, almost one-third of the installations were performed by six companies, half of which were in the auto industry (Ayres & Miller, 1981). The technology did not yet permeate the whole manufacturing industry.

These historical developments, including the decline in manufacturing employment, the increase in robot and language production, and the gap with Japan in robot adoption, characterized the broader context in which the Robotics Institute began to operate. The government's concern summarized by Cyert as "the Japanese are eating our lunch with robotics" to Reddy, Traub, and Newell was a verbal snapshot of the developments underway, which sowed the seeds of the Robotics Institute at CMU. The newly launched institute, which positioned robotics as an entrepreneurial project with its partners and funding, thus took on a larger mission than merely advancing the state of the art. Attached to the technological aspirations at the Robotics Institute was the aspiration to help salvage the United States' position in international trade via a robot-enabled increase in productivity, product diversity, and exports.

The traces of this aspiration can be found in a 1981 report in the archives of the Institute, emphasizing the potential of computer control and reprogrammability of robots for this rescue mission. Written by Robert Ayres (1981), a professor of Engineering and Public Policy, with his Ph.D. student Steven Miller (1981), this report identified where in the US industrial production and economy computer-controlled reprogrammable robots could boost the value created by the industry, and reclaim the country's competitive edge. According to the report, job shops, batch production, and mass production were the production methods employed in the 1980s (Ayres & Miller, 1981) (Figure 5). Job shops would demand higher labor skills and general equipment, which would enable the production of a variety of products but in small quantities at a time (Ayres & Miller, 1981). Conversely, mass production would not require the same level of skill and would necessitate specialized equipment to produce one product in much larger quantities (Ayres & Miller, 1981). This was a well-known problem in the industry, called the "tradeoff between flexibility and efficiency," as Steven Miller told me forty years later in an interview. The efficiency of machines in mass production, namely, hard automation systems, would decrease the cost per unit of a product, but these systems had to be reconfigured at an impractically high cost to mass-produce a new product.[26] It was cheaper to retire them than to rebuild them.

Based on this evaluation, Ayres and Miller (1981) presented batch production as the sweet spot for efficiency and flexibility, involving mass production for shorter periods and potentially with reconfigurable machinery to produce a variety of products in batch-size quantities. According to the report, a large share of manufactured goods and of the total value generated by durable goods in the

24 Stokes (1982) quotes Paul H. Aron, executive vice president of Daiwa Securities America Inc., who informed him about "11,250 programmable robots at work in Japan," which was "almost three times the number in the United States."
25 Stokes (1982) quotes Paul H. Aron's projection.
26 Steven Miller explained this during his interview with me in 2021.

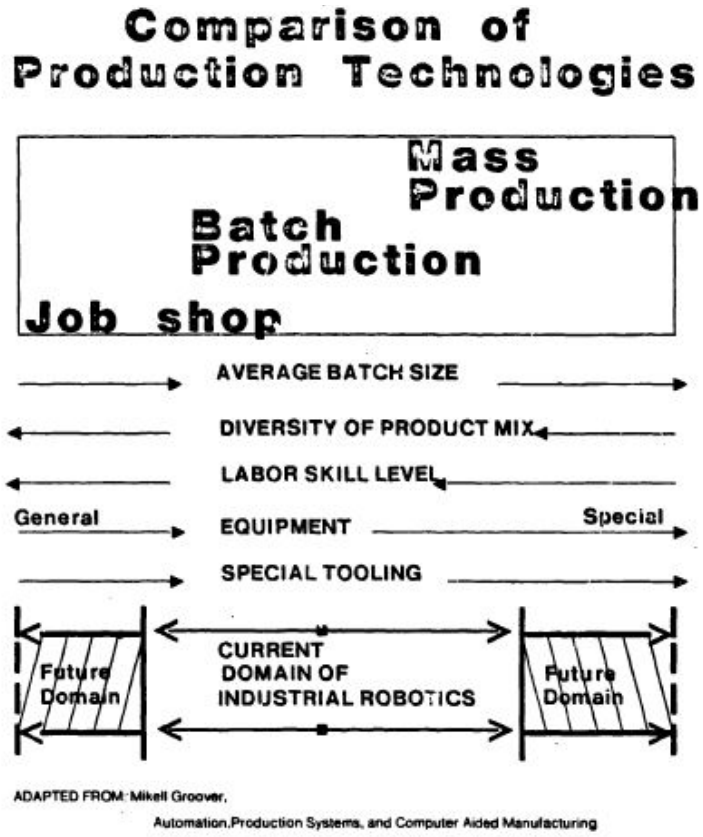


Figure 5. Diagram of production technologies in the 1980s (Ayres & Miller, 1981) (Source: Steven Miller)

US had already been relying on this production method (Ayres & Miller, 1981). However, with a low percentage of reprogrammable machine use, the batch production of these goods was predominantly based on manual and labor-intensive operations (Ayres & Miller, 1981). Ayres and Miller thus showed that the manufacturing industry had been doing the right thing with the wrong tools. Computer-controlled reprogrammable robots at the time were not as fast as hard automation systems with fixed sequences, but they were the perfect tools for expanding programmable automation in batch production. Their adaptability for new sequences held promise for the cost-efficient production of a variety of products. On the CMU campus, these concepts would soon find their place in a similar effort envisioning programmable automation but for a different industry with a very distinct way of producing: building construction.

2.3.2 Machines Operating Machines

Among industry actors, Tom Murrin and the Westinghouse Electric Corporation were also aware of the promise of productivity increase through programmable automation. The company entered the robot market in 1981, and only a year later, it acquired Devol and Engelberger's Unimation, founded after Unimate, and previously developed PUMA and the robot programming language ("Sale," 1982). In addition to this acquisition, Westinghouse formed strategic partnerships—ironically—with the Japanese companies Kawasaki, Komatsu, and Mitsubishi to quickly grow its product line based on the licenses held by these companies (Shim & Kumar, 2005). The Robotics Institute was the university research arm of this ambitious effort, and the Institute's goal, in Raj Reddy's (1989) words,

was to help its sponsor “leap-frog competition” in manufacturing. Based on this premise, between 1980 and 1984, the Robotics Institute developed a list of applications in collaboration with Westinghouse, including a constraint-directed search system for scheduling, a flexible manufacturing cell for machining turbine blades, and an automatic visual inspection system for inspecting printed circuit boards.[27]

Among these systems, the Flexible Manufacturing Cell stood out by offering a full display of the key advances in the 1960s and 1970s I described earlier in this chapter. Built at Westinghouse’s Winston-Salem plant, this cell was a salient demonstration of programmable automation, in which two computer vision systems and two industrial arms were programmed with a high-level language to operate an open die forging workflow (Figure 6). As the two developers of the system, David Bourne and Mark Fox (1983), describe, the workflow starts with the first vision system detecting a billet to be processed.[28] Once the detection is done, the first arm transfers it into a rotary furnace for heating and from there into the open die forge beating the hot metal to form an elongated blade shape. Then, the second arm moves the blade to a vision-based inspection that reconstructs the item’s three-dimensional model. After reconstruction, the second arm transfers the blade to a cropper that shaves the material residue on the item, then to a stamper that puts a batch code on the item, and finally to a rack for the item to be transferred to the next cell. A VAX 11/750 superminicomputer and a rule-based control system named TRANSCELL developed by the researchers (1983) control the whole automation system.

Bourne and Fox (1983) designed TRANSCELL to handle high-level user input, database management, and communication between the arms and the existing machining equipment at three levels (Figure 6).[29] At level one is the language interpreter through which an operator could write the problem into the system’s high-level language as a “rule”—essentially “a set of actions” with “a set of preconditions”—and enter it into the system (Bourne & Fox, 1983; Bourne & Fussell, 1982). An example problem would be “lower[ing] the furnace’s temperature during a long vacation,” as Bourne and Fox (1983) describe in their introduction of the system. At level two is the database that receives the rule input from the interpreter and the data—indicating the states of the machines—from the controllers of those machines at level three. In the furnace example, the new rule input updates the expected temperature data, the database accordingly sends a command to the controller of the furnace to execute the change, and the controller sends the new temperature data back to the database. The system controls the arms in the same way. The database sends commands to the arms’ controllers to execute their motion.

‘Aside from technical knowledge, the Flexible Manufacturing Cell/TRANSCELL embodied the values underpinning the CMU-RI’s foundation and the factors in the broader context in which the Institute operated. This system reflected the Institute’s entrepreneurial spirit and inclination to build systems for the real world, which stemmed partially from Cyert’s managerial style of administration. The system was a combination of technologies, which became discursive artifacts in the discussion on the United States’ declining manufacturing industry and its critical position in the global economic race against Japan. In a nutshell, a control system allowing high-level programming input was developed to customize arms and demonstrate their task versatility in a Westinghouse

27 See more on the technical details of these systems in the *Robotics Institute Annual Research Review*: (1) see the ISIS system, an AI-based constraint-directed scheduling system, in David Bourne and Mark Fox’s (1983) article “Autonomous Manufacturing: Automating the Job Shop” published in the 1983 Review, (2) see the Flexible Manufacturing Cell in the same article as well as in Paul K. Wright’s (1983) “A Manufacturing Hand” also published in the 1983 Review, and (3) see the Automatic Visual Inspection system in Robert Thibadeau’s (1984) “Automatic Visual Inspection as Skilled Perception” published in the 1984 Review.

28 The description of the system in this paragraph is based on Bourne and Fox (1983).

29 The description of system components in this paragraph is based on Bourne and Fox (1983) unless stated otherwise.

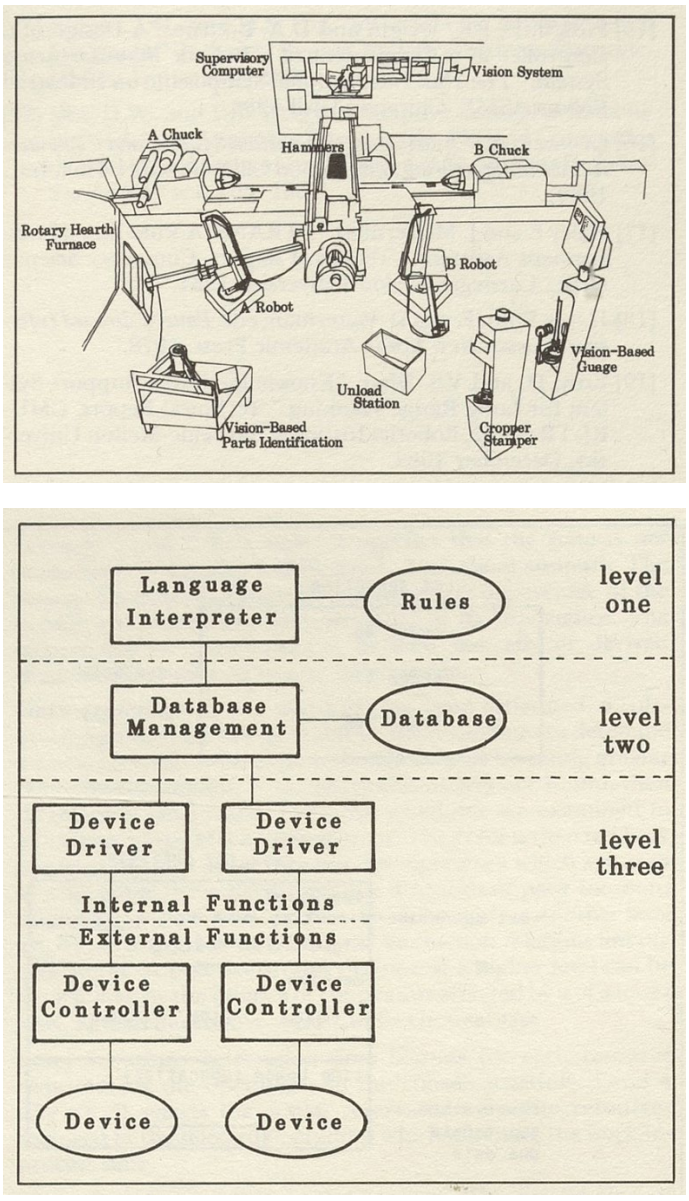


Figure 6. An illustration of the Flexible Manufacturing Cell (top) and the system diagram of TRANSCELL (bottom) (Bourne & Fox, 1983) (Source: *The Robotics Institute 1983 Annual Research Review*).

manufacturing cell. The system was also a product of the aspiration to revive the US manufacturing industry, made possible by the Westinghouse Corporation's funding and facilities.

A combination of factors—globally, the proliferation of robots and high-level languages and the high hopes for the country's manufacturing industry built around these technologies, and locally, the Westinghouse partnership—generated a robot-productivity discourse at the CMU-RI. Transitioning into programmable automation via computer-controlled reprogrammable arms was prescribed as a way to boost the US economy through higher batch production output at a lower cost and with less labor, as stated in the Ayres and Miller's (1981) report. Technologies that could facilitate this transition were produced hand in hand with Westinghouse. "Unmanned," as it would sometimes be described in articles, the Flexible Manufacturing Cell was one of those, which I showed closely.

As I briefly noted before and illustrate further in the coming sections, a similar discourse appeared as the ideological underpinning of construction robotics during its formative years at CMU's DCE in the early 1980s. It conditioned the broader vision of the civil engineers kickstarting this interdisciplinary field and became their default motivation in research. From a high-level perspective, robotic technologies for productivity seemed to be an unquestionably good proposition applicable to both manufacturing and construction in the early 1980s.

2.4 Situating Robotics in Construction: Robotics at the Department of Civil Engineering

2.4.1 "Mobile Work Object" Versus "Stationary Work Object"

Fast forward to 1989, the civil engineering and robotics community in the United States was at a moment of reflection on the robotization of construction. The reflection was on the challenges posed by construction work to robots and the missing capability, mobility, needed to fit these technologies to dynamically changing settings on job sites.

"Researchers at the Robotics Institute of Carnegie Mellon University are trying to make the transition to robotics a smooth one through work being done in the Civil Engineering and Construction Robotics Laboratory," states the beginning of a book chapter by Richard Miller (1989c) reporting on the field. After listing the robots' benefits as "decreas[ing] the number of hazards," "perform[ing] superhuman and unhuman tasks," and "increas[ing] the productivity rate," Miller (1989c) asserts that construction tasks would "demand" different robot "capabilities" than working "in a fixed position and performing a highly repetitive task in an assembly line format." Miller (1989c) follows with a list of these capabilities, and, unsurprisingly, "mobility" is at the top, succeeded by "the ability to work with large forces and weights" and "objective sensing in various environments." These capabilities were considered the baseline to make robots handle construction work in the chapter (Miller, 1989c). In reality, they had only been partially achieved by 1989, and the idea of robots on job sites was thus still far off for US construction.

Published in the same year, an article reviewing construction robot prototypes in the US went deeper in comparing construction and manufacturing work. Written by Miroslaw Skibniewski (1989), a professor of civil engineering at Purdue, and one of his post-doctoral researchers, Jeffrey Russell (1989), the comparison reads as follows:

> *Unlike batch manufacturing, where the work object is mobile through the production facility and work tools can be stationary, the "work object" in construction is stationary, of large dimensions, and constantly changing as work progresses. Tools are mobile, whether hand-held or mechanized, and are operated by humans working on a stationary, massive, and changing object.*

Before Purdue, Skibniewski was just a walk away from the Robotics Institute doing his Ph.D. at the DCE. He was likely knowledgeable about the CMU–Westinghouse projects, which were featured in the Institute's research reviews every year. His succinct description above clearly traces a key quality of construction work—"mobile tools" on a "stationary work object"—and foregrounds it as a challenge in industrial robot adoption on sites. Stationary arms, in CMU's Flexible Manufacturing Cell, for instance, could conveniently move small-size billets from one station to another. But in construction, it is the arm that must be moved on an unstructured site to apply changes to much larger building parts. Movement on an unstructured site demanded more advanced robot perception and navigation than the state-of-the-art machines had in 1989.

2.4.2 Perception, Navigation, and the High Hopes for Construction

At the beginning of the 1980s, the idea of dealing with a stationary work object with a moving robot was an even more distant vision. As a result of the developments in the 1960s and the 1970s I discussed earlier in this chapter, robotic arms were already mature at the time. Conversely, robot perception and navigation research was behind, and thus, mobility was still the missing piece towards building robotic tools for on-site construction in the United States. In the early 1980s, this began to change with two mobile robot prototypes developed by the Civil Engineering and Construction Robotics Laboratory (CECRL) based in the CMU DCE.

As I describe in this section, I see two significant effects that these prototypes had in drawing construction and robotics closer at the DCE. The first one is a particular approach to technology design that was at play at the Department, which focuses on building prototypes and extending machine capabilities to fit them into construction. At the turn of the 1980s, this was a distinct move towards transitioning from generic mobile systems to systems specifically designed for construction. The second effect was, other than a sense of awe, seeing these robots in action sparked hope for the robotization of tasks in construction at the DCE, similar to what arms and high-level programming languages had done for manufacturing as I discussed earlier in this chapter.

At the verge of the Robotics Institute's launch, state-of-the-art mobile systems were not yet fully suited to outdoor environments, described in the field of robotics as unstructured and dynamically changing. They were not designed for any specific application domain or site conditions. The renowned systems were the John Hopkins Beast by the Adaptive Machines Group, the Stanford Cart by SAIL, and Shakey by SRI, which were all customized hardware platforms built to conduct sensing and path-planning research.[30] These prototypes were only able to navigate relatively well-structured spaces such as rooms or corridors. For instance, even Shakey, as the pinnacle of these systems equipped with a TV camera, an optical rangefinder, and "cat-whisker" touch sensors, was typically test-run in contrived indoor setups (Nilsson, 1984) (Figure 7). Perception and navigation research via such prototypes in the US was concerned with sensor integration and algorithms for mapping, path planning, and obstacle avoidance. However, the outcomes were understandably far from the level that would make robots fit for construction.

In the initial years of the CMU-RI, perception and navigation research was primarily carried out by a group called the Mobile Robot Laboratory, which was started by a researcher named Hans Moravec. Like Raj Reddy, Moravec was a Ph.D. student at SAIL, known for the vision system he developed for the lab's famous mobile robot, the Stanford Cart.[31] Relying on Moravec's SAIL experience and ONR sponsorships, the Mobile Robot Lab mainly focused on autonomous navigation research and systems to be deployed on underwater robots (The Mobile Robot Laboratory Staff, 1984). The research was conducted on three prototypes the lab (1984) built as testbeds and named after planets.[32] Pluto was the first one, which the group used to develop a hardware and control system for omnidirectional drive. Neptune was the second one, through which the group developed a stereo-based and sonar-based obstacle-avoiding system and a path-planning algorithm. The third one,

30 See more on these robots in various sources: (1) The John Hopkins Beast in a video demo titled "APL Flashback: 1965 – The "Hopkins Beast" Mobile Automaton" on JHU Applied Physics Laboratory Youtube channel (2012), (2) The Stanford Cart in Hans Moravec's (1983) article titled "The Stanford Cart and the CMU Rover" as well as a video demo of the robot in the New York Times (2010) available online, and (3) Shakey the Robot in an SRI report edited by Nils Nilsson (1984) as well as a video demo titled "Shakey the Robot: The First Robot to Embody Artificial Intelligence" in SRI International Youtube channel (2017).

31 See Moravec's (1980) dissertation titled "Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover."

32 See more on the lab's prototypes in an overview by the Mobile Robot Laboratory Staff (1984) within the *Robotics Institute 1984 Annual Research Review*. The brief description of these prototypes in this paragraph is based on that overview.



Figure 7. Shakey in a lab setup (Source: SRI International)

Uranus, was used for testing omnidirectional drive with specialized wheels. Despite these significant technological advancements, these three were hardware platforms that mainly traversed contrived test set ups similar to their counterparts in the previous two decades.

Perhaps due to the civil engineering knowledge of its members, the CECRL's prototypes in the early 1980s were shaped by a different approach, which connected civil engineering and construction as a work domain with mobile robot research. That was to build ready-to-deploy vehicles for outdoor civil engineering and construction tasks, more broadly and in its founder William Whittaker's words, as "intelligent machines at work in the world" rather than platforms for lab research.[33] Alternative to building generic systems and conceiving them as applicable to several application domains, this approach concerned conceiving and building systems specifically for heavy civil engineering work, such as nuclear cleaning, mining, or construction. The lab sought to ensure the fitness of prototypes for specific work conditions. This was a major push towards building machines for unstructured, dynamically changing sites.

The two important outcomes of this approach was the Remote Reconnaissance Vehicle (RRV) and the Terregator. The Remote Reconnaissance Vehicle (RRV) was a response to a nuclear accident in 1979 at Middletown Pennsylvania, which provided an opportunity and a modest grant for the lab to start developing it (Carnegie Mellon University, 2009).[34] "Somebody else's trouble," as Whittaker refers to this infamous Three Mile Island accident in an interview with me four decades later, "was somebody's silver lining." Ready to be deployed in 1984, the RRV was an unmanned ground vehicle equipped with multiple cameras, lights, and radiation monitoring equipment (Figure 8).[35] The RRV was soon followed by the Terregator, a prototype with basic self-drive capabilities also developed contextually with mining and construction tasks in mind.

33 William Whittaker used this phrase during our interview in 2021. As he also described to me in the interview, for this mission, he trained himself in electrical engineering, mechanical engineering, and computing since his undergraduate years at Princeton University.

34 Based on my interview with Whittaker in 2021 and a short article titled "Remembering Three Mile Island" published on Carnegie Mellon University website (2009). Whittaker's group that developed the RRV included Jon Bares, Jim Osborn, Leona Champeny Bares, Brian Smith, Chris Fromme, Gary Baun, and Clark McDonald (Carnegie Mellon University, 2009).

35 Based on a video about the robot on CMU Robotics Institute Youtube channel (2019).

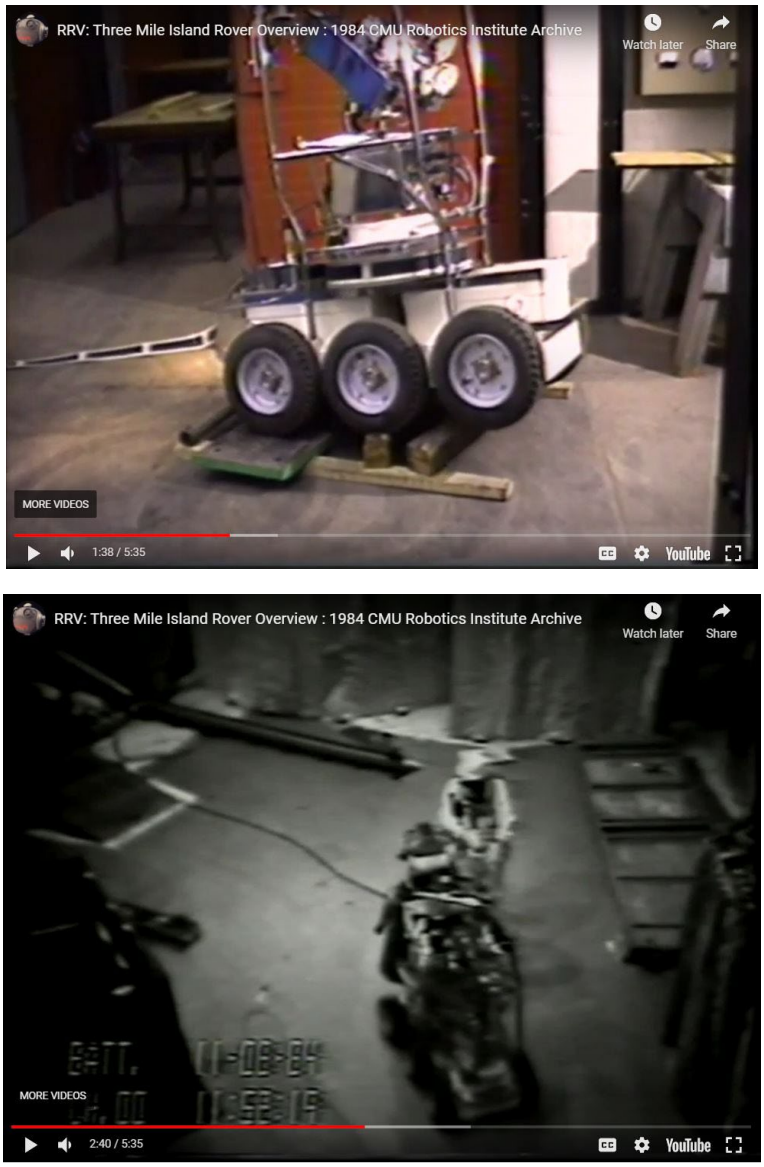


Figure 8. The RRV in action (Source: Carnegie Mellon University Robotics Institute, Field Robotics Center)

Real-world context and constraints were important considerations in designing and engineering the two systems. Along with focusing on sensors and algorithms, the group was rigorously engineering the final physical form of their prototypes to withstand the outdoors. The RRV was assembled by taking the site conditions into account. With its six-wheeled drive base, moving on rugged surfaces, crossing over small obstacles, and climbing slopes were some of this robot's extensively tested and established capabilities.[36] Similar to the RRV, the Terregator featured a six-wheeled drive base, but it had improved locomotion, allowing it to navigate along intended paths and in more challenging conditions, such as natural terrain, stairs, and slopes.[37] The two prototypes were also cornerstones of perception and navigation research, demonstrating the beginnings of reliable teleoperation and sensing as well as preliminary self-drive capabilities. The RRV's capabilities were demonstrated as the vehicle was teleoperated on the basement of the Three Mile Island nuclear plant, surveying and

36 This short description of the RRV is based on the video on CMU Robotics Institute Youtube channel (2019) noted earlier.
37 This short description of the Terregator is based on a video on CMU Robotics Institute Youtube channel (2016).

detecting radiation.[38] The Terregator was showcased in an awe-inspiring demo at the southern edge of the CMU campus between Schenley Park and Porter Hall, which houses the DCE to this day.[39] Along with the RRV's Three Mile Island deployment, this famous demo, in which the vehicle climbed up the Schenley Park terrains and gently rolled down the Porter Hall entrance stairs, showed the possibilities offered by robots taking their first steps from well-curated unstructured lab setups into unstructured and dynamically changing settings (Figure 9).[40]

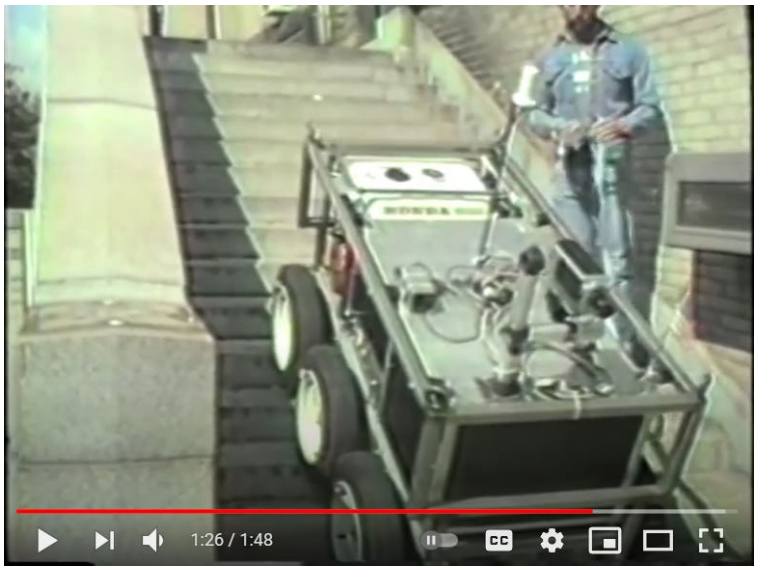


Figure 9. The Terregator during the demo on CMU campus (Source: Carnegie Mellon University Robotics Institute, Field Robotics Center)

Seeing the RRV and the Terregator navigate in the wild fueled the imagination of robots as construction tools at the DCE. As faculty members Irving Oppenheim and Chris Hendrickson reflected after four decades, the CECRL and the two prototypes reconditioned their view and aspirations on the robotization of construction.[41] The prototypes demonstrated to the CMU-RI and the DCE that the idea of systems combining manipulation and navigation and operating on sites might not be too far off. They brought hope and excitement that "mobile tools working on stationary, massive, and changing object[s]" in Skibniewski and Russell's (1989) later description, may turn from a distant vision into a somewhat promising conception.[42]

2.4.3 Same Technology, Similar Discourse, Different Work

Alongside the momentum the Robotics Institute created in the early 1980s, robotics also began to grow as a research stream in CMU's DCE. The early adoption of robotic arms in manufacturing painted a picture of success for a group of civil engineers at the Department. In addition to that, the

38 Based on my interview with William Whittaker in 2021.
39 Based on my interview with Irving Oppenheim in 2021.
40 These two prototypes in the wild put William Whittaker in the spotlight, and convinced Raj Reddy to channel resources for him to continue as part of the Robotics Institute. William Whittaker told his encounter with Raj Reddy and his transition to the Robotics Institute during his interview with me in 2021. Whittaker left the DCE with the Terregator, which became the platform for the ongoing autonomous navigation research at the Institute. See the overview by the Mobile Robot Laboratory Staff (1984) within the *Robotics Institute 1984 Annual Research Review* noted earlier.
41 Based on the engineers' respective interviews with me in 2021.
42 The phrase, "mobile tools working on stationary, massive, and changing object[s]," is a mildly edited quotation from Skibniewski and Russell's (1989) article.

RRV and the Terregator, demonstrating preliminary self-drive out in the field in unstructured settings, showed them that the idea of arms becoming mobile construction tools got a step closer to materialization. Not far from the Institute aspiring to help the transition of the US manufacturing industry with robots, a group of civil engineers at the Department began to envision a similar transition in the nation's construction industry. Fueled by the excitement around arms as well as the two mobile robots, they were building an interdisciplinary field engaging civil engineering, construction, and robotics.

In this section, I first describe the forces at play behind the emergence of this field: knowledge transactions with the Robotics Institute, the excitement created by the CECRL and the lab's two mobile robot prototypes, and the ambitions to extend robotics research into on-site building construction in the United States. I reveal that this pioneering sociotechnical project produced, alongside technical research, aspirations to reconfigure construction work around robotic technologies. Looking into the civil engineers' discourse carrying these aspirations, I trace a couple of salient elements. The first is their emphasis on productivity for the US construction industry, often viewed as behind in technology adoption next to other industries. The second is the engineers' descriptions of construction domain depicting a high-level view of it. These descriptions highlight the domain's distinctive aspects from those of manufacturing to point out the challenges toward robotization and its similarities with manufacturing to make a case for robotization.

Ultimately, I point out that the civil engineers at the DCE situated the idea of intelligent machines and a productivity discourse resembling the one built at the CMU-RI around these technologies within the construction domain. However, as I discuss in this section, that particular domain required a drastic shift from hand tools to intelligent machines compared to the one in manufacturing that entails intelligent machines operating the existing. Further, while highlighting the differences shows that the engineers already recognized the complex makeup of workplaces in construction than those of manufacturing, highlighting the similarities shows they conceptualized part of on-site construction work as more mechanical and analogous to manufacturing work than it really is. Lastly, the engineers did not incorporate significantly in their research the contingency of a set of building tasks and workflows on materials critical to inform and evaluate technologies. They weighed their focus more on advancing autonomy, sensing, or motion needed to operate machines on sites.

An important force that drew construction and robotics closer at the DCE was the influence of the Robotics Institute, which became a vibrant center connecting faculty and students across campus. A couple of outliers at the DCE were among the early ones to engage with the Institute and conduct research outside the Department's main focus (computer-aided design expert systems). Irving Oppenheim, who was at the Department since the 1970s, was one of them. After four decades, he excitedly remembered Marc Raibert—who later founded Boston Dynamics—and his robot locomotion reading group, which he "religiously" attended.[43] The other outlier was Chris Hendrickson, who joined the Department in 1978. "As a researcher, you got to have kind of a sense of where things are happening," said Hendrickson as he explained to me how his research got extended into robotics.[44] Seminars at the Institute were Hendrickson's resources while he had to "continually learn new techniques" such as "neural nets" to operate in robotics.[45]

The other force came internally from Whittaker's CECRL. By raising excitement around self-drive with the RRV and the Terregator, this modest initiative with little funding created hope for the

43 Irving Oppenheim mentioned about Marc Raibert's reading group in my interview with him in 2021.
44 Based on my interview with Chris Hendrickson in 2021.
45 Ibid.

robotization of construction tasks through mobile manipulators. The laboratory and the prototypes expanded the imagination of robots as building instruments and encouraged other robotics activities to start at the DCE. They created a strong synergy around robotics, which Whittaker characterized as "the new character on the block" in the context of engineering.[46] "It was up and coming, new, fresh, and pretty exciting," said Whittaker in our conversation, helping me understand the growing interest in robotics for various engineering applications. "That's the basis on which I was able to bring it [robotics] to life and attract the teams, the community, and the work," he continued, describing his individual efforts to bring the two fields together at the Department. His casual weekly seminars were attended by students and, at times, other faculty members to discuss early ideas for the automation of excavation or building trades, some of which the lab would then pursue.[47]

The other force were the efforts by a list of few key figures, inspired by the robotics research for on-site building construction in Japan, and aspiring to establish a similar research agenda at the DCE. The list starts with Dwight Sangrey, who was the Department's head between 1979 and 1985. Sangrey's time at CMU coincided with the boom in corporate-led robotics research in the Japanese construction industry, resulting in big contractors, such as Shimizu and Kajima, developing and deploying single-task robots on sites.[48] Sangrey, who knew of these advances, initiated an agenda with a focus on panelized construction at the Department.[49] Sangrey also contributed to the institutionalization of construction robotics research globally by organizing two international workshop-style events at CMU in 1984 and 1985, which later continued as ISARC (IAARC, 2021). Another key figure was Abraham Warszawski, a collaborator of Sangrey invited from Technion to help move the panelized construction agenda forward.[50] During Warszawski's stint at CMU, the two engineers mainly focused on the analysis of construction work for robot development and adoption. The last figure concentrating on on-site building construction was a Ph.D. student of Hendrickson's named Miroslaw Skibniewski. Also knowledgeable of Japanese construction robots, Skibniewski was researching the economic feasibility of robot applications to surface tasks in his Ph.D.[51]

While knowledge transactions with the Institute, the excitement around the CECRL's mobile robots, and the development of the new research agenda set the technological underpinnings of construction robotics, the ideological underpinnings of this emerging field were also in the making at the DCE in the 1980s. The civil engineers were situating robotics in construction technically, but perhaps more interestingly, through discursive means. A closer look at their discourse in publications unravels their aspirations to reconfigure construction work around robotic technologies. These aspirations were weaved based on a couple of characteristic elements.

The first one is an emphasis on productivity in construction. The engineers often employed an evidence-based narrative, offering an outlook on the productivity decline in the US construction industry relative to its scale and importance for the nation. Publications often began with a list of key

46 Based on my interview with William Whittaker in 2021.
47 Ibid.
48 As researchers Thomas Bock and Thomas Linner (2016) point out, this strand of single-task robot research and development in Japan represented a transition from "the shifting of complexity [of construction] into a structured environment" to "the development and deployment of systems that could be used locally on the construction site." See the various resources to learn more about these efforts: (1) Tomohiro Ueno and colleagues' (1986) conference proceeding titled "Construction Robots for Site Automation" that overviews research by Shimizu Corporation, (2) Miroslaw Skibniewski and Jeffrey Russell's (1989) "Robotic Applications to Construction" for a list of construction robots developed in Shimizu, Kajima, and Ohbayashi by 1989, (3) Bock and Linner's (2016) book *Construction Robots: Elementary Technologies and Single-task Construction Robots* that includes a survey of single-task robots developed in various Japanese companies.
49 This information is based on my interview with Irving Oppenheim in 2021.
50 Ibid.
51 Miroslaw Skibniewski's (1986) Ph.D. dissertation is titled "Engineering and Economic Analysis of Robotics Application Potential in Selected Construction Operations." Part of the dissertation is also discussed further in this chapter.

numbers calling out the need for improvement. Similar to the robot productivity discourse built on hopes for the manufacturing industry and motivated the systems at the Robotics Institute, boosting the productivity in the construction industry larger than manufacturing at the time became the answer to why robots' use in construction mattered. An example is William Whittaker's (1986) proceeding for the CAD and Robotics Conference, where he outlined the motivation for using robots by emphasizing the scale of the industry, "representing six to ten percent of the United States Gross National Product," and contrasting it with its "productivity," which had been "in decline for two decades."[52] Whittaker (1986) accompanied this item on the list of motivations with "quality assurance," "better control over the construction site," and the crucial aspect of worker "health and safety." A year before, Abraham Warszawski and Dwight Sangrey (1985) conveyed the same industry size–productivity contrast with different numbers: "8% of the Gross National Product," "40% of the total private domestic investment," and "1.5% productivity decline per year."[53] By stating that the construction industry was too big to let its productivity decline, this fact-based narrative was making a case for robots.[54]

By extension, civil engineers often associate robots with productivity assurance in construction in publications. Their statements at times exuded from case studies or estimations. However, for the most part, they were hopeful presumptions, potentially or partially right but also uncontested in the spirit of the 1980s bearing excitement. An example case study was performed by Miroslaw Skibniewski and co-authors (1986) on the "cost and design impact" of robots in surface tasks.[55] Published within the 1986 CAD and Robotics proceedings, the study presented an analytical argument on productivity at the task level, stating that "worker productivity [in surface tasks] decreas[ed] with time due to tiredness" (Skibniewski et al., 1986). Robots could bring "significant benefits" in this area, the study advocated, supporting its argument based on Shimizu's SSR-2 fireproofing robot, which cut down the execution time on a "work unit" from "51 minutes" to "about 22" (Skibniewski et al., 1986). On the other side, it was common to see projections or early conclusions, although no robots were used or evaluated in the US construction industry. "There is widespread interest in applying robotics to construction, to bring productivity gains to this large but diffuse industry …" Oppenheim and Skibniewski (1988) wrote in the International Encyclopedia of Robotics. "The potential for using robotics in construction is very great with associated benefits to owners, builders, and labor" was a similar projection by Warszawski and Sangrey (1985).

However, the same technology that had previously shown results in the manufacturing industry would not directly equate to the same productivity levels promised by the engineers in the construction industry. Although arms and mobile robots seemed to be a good proposition for construction, the difference in automation levels between the two industries was a caveat that this proposition overlooked. Manufacturing in the 1980s was automated to a degree, so robotic systems, such as CMU's Flexible Manufacturing Cell I previously discussed in this chapter, represented a transition from hard to programmable automation. The technical explanation for this was that these systems were being layered onto cells with existing machinery, and in most cases, it was not the

52 This conference was the third International Symposium on Automation and Robotics in Construction (ISARC) and the first symposium version of ISARC held at Marseilles France in 1986, following the first two conducted in workshop format at Carnegie Mellon in 1984 and 1985. William Whittaker's (1986) proceeding that included this point he made is titled "Construction Robotics: A Perspective." Whittaker's (1986) point on productivity decline is based on Skibniewski's (1984) dissertation proposal titled "Strategy for Engineering and Economic Analysis of Robotics Applications Potential in Selected Construction Work Operations."

53 Warszawski and Sangrey (1985) mentioned these numbers in their article titled "Robotics in Building Construction," also mentioned later in this chapter. The engineers (1985) referenced Hans-Jürgen Warnecke and Rolf-Dieter Schraft's (1982) book titled *Industrial Robots: Application Experience* for a part of these numbers.

54 Here I borrowed and tweaked the term "too big to fail" in economics. According to Julie Young's (2023) short article "Too Big to Fail: Definition, History, and Reforms" on Investopedia website (2023), the term refers to "a business or business sector so ingrained in a financial system or economy that its failure would be disastrous."

55 This publication is based on part of Skibniewski's (1986) Ph.D. dissertation previously mentioned.

robots but these tools that performed the core machining tasks.[56] In simple terms, it was the contemporary intelligent machines that operated the conventional ones in these setups. However, for the construction industry, robots would essentially represent the first foray into automation. The expectation was more of a leap from handheld and power tools to intelligent machines than a moderate transition. In that sense, robots' relative success in manufacturing was not necessarily a guaranteed model for construction.

Other than the emphasis on productivity, high-level descriptions of construction work domain, comparing or contrasting its aspects with those of manufacturing work domain were used to present aspirations. Alluding to the challenges toward robotization by designating the difference between the typical sites of these two work domains was common in publications. "Dynamic," "unpredictable," and "unstructured" were the keywords the civil engineers used to distinguish job sites from factory floors in their often brief descriptions to foreground challenges to be resolved to achieve robotization. For instance, Whittaker (1986) identified "initiating tasks in unstructured environments to accomplish goals" as the essence of construction work, and compared those settings to those in manufacturing: "Structured environments, like those found in factory settings, do not admit the dynamism and uncertainty attendant with unstructured environments." Similarly, Warszawski and Sangrey (1985) highlighted the "ever changing tasks and an unstructured workplace" in construction, and for that reason, they concluded that "the fixed sequence mode of operation has a very limited application in construction." "The volatile and unpredictable nature of the construction environment" was also the reason behind the "tendency to use old and proven solution instead of innovative methods and technologies" in the industry according to Warszawski and Sangrey (1985).

The engineers' remarks were not wrong. The level of organization on construction sites was nothing like that in the factories of the 1980s, where work was organized sequentially through cells that concentrated specialized machines and human labor. As I previously noted with Skibniewski and Russell's (1989) depiction in this chapter, in factories, the job was about moving appropriately sized parts between these machines and workers, whereas on construction sites, it was about moving the machines and workers to build large parts. The civil engineers were already aware that construction sites would be more demanding and challenging for robot adoption.

Relatedly, the words "dynamic" and "unstructured" contained a particular difference of construction work that the civil engineers identified and used as the basis of a technological prescription. That is, construction work is more of a spatial practice than manufacturing. "Spatial" means that the work entails positioning or applying changes to an often large-scale, complex-shaped, and, at times, unorderly positioned object among a plethora of others in space. Therefore, it demands a higher level of mobility between points and perception of these objects. Simply put, for a construction worker, how an object is positioned in space, what surrounds it, and how to navigate among these objects matter while performing tasks. In the 1980s, the contrived layout and control of objects at factories was not easy to engineer on construction sites. Based on that, for instance, the engineers knew robot perception and navigation, which were at their infancy in the 1980s, had to advance in order to make robotic arms mobile and more functional in on-site construction.

Interestingly, some high-level descriptions foregrounded part of construction work's resemblance to manufacturing for making a case for robots. Similar to the engineers' description of differences, these descriptions of resemblance were also limited to a few high-order words such as "arduous" or "repetitive." Because of these shared qualities, certain building tasks could be as suitable for

56 See more in Richard Miller's (1989b) "The Robot Manufacturing Cell" chapter in his book *Industrial Robot Handbook: Case Histories of Effective Robot Use in 70 Industries* previously mentioned.

robotization as manufacturing tasks as the engineers were inclined to believe. For Miroslaw Skibniewski and coauthors (1986), these were surface finishing tasks, such as coating, painting, and sandblasting, not only because robots in these tasks would potentially bring "rapid economic payoff," but also the work itself "closely resembles some of the arduous, repetitive work tasks already automated or roboticized in the manufacturing industries." Similarly, Abraham Warszawski and Dwight Sangrey (1985) also referred to a set of tasks including painting and concreting, and identified them as having "homogenous and well-structured portions," which "preprogram robots" could be suitable for.

What we see in all these high-level descriptions is that the civil engineers differentiated unstructured job sites from more structured factories, but some also conceptualized part of on-site construction work as more mechanical and analogous to manufacturing work than it really is. More specifically, they take certain building tasks as closely repetitive as manufacturing tasks. What went unchecked with these descriptions is that repetition in manufacturing tasks with the existing machines occur in a different resolution than in construction tasks. In the 1980s, these existing technologies were already central to organizing manufacturing work in almost perfectly, tightly planned, and repeated cycles of discrete action. This organizing element in construction was—and still is—the building trade system, and the varying order of workers' actions on sites. Most trades' work occur in loosely planned, broader repeating cycles of a larger mix of actions and numerous improvisations. In that sense, the manufacturing industry had an existing layer of mechanization establishing this fine-grained repetition, which enabled the relatively successful integration of arms as inspectors, movers, or operators into existing cells from the 1960s onwards. The robotic arms in CMU's Flexible Manufacturing Cell integrated with the existing open die forge at the Westinghouse plant was an example. Conversely, with its lack of existing mechanization, the construction industry was not yet ready for a new layer of order or the integration of mobile arms as builders on sites.

Lastly, these high-level descriptions missed out on the material-oriented character that certain construction tasks have, which is critical to take into account in researching and developing robotic technologies. "Material-oriented" means that the work entails installing a building unit, such as a brick, in a series of precise locations or applying a chemical substance, such as paint, to a series of surfaces. The work involves human action, but, more importantly, an outcome of a building part that should meet certain quality expectations and technical standards. How that outcome meets the standard depends not only on the workers' or a machine's actions but also on how workers or machines work with those materials. Although some material-oriented tasks, such as painting, plastering, or sandblasting, were featured and examined within the work of Warszawski and Sangrey (1985), and Skibniewski (1986), the weight at the DCE in the 1980s was mainly on advancing machine autonomy, sensing, or motion needed in tasks.

This may perhaps be due to the excitement around AI and robots that the Robotics Institute created as well as the limited hands-on work on sites and with technologies back then at the DCE. Other than the CECRL's mobile robots, one may also see Irving Oppenheim and Chris Hendrickson's projects as an outcome of that excitement. Inspired by Raibert's robot locomotion reading group at the Institute, Oppenheim redirected his own research towards dynamically stable systems based on his Ph.D. student Lyman Petrosky's work and a Department of Energy grant.[57] He co-published a series of work on construction robotics focusing on autonomy or motion until the mid-1990s.[58]

57 Based on my interview with Oppenheim in 2021.

58 See Lyman Petrosky and Irving Oppenheim's (1988, 1989) publications titled "Application of Model Based Control to Robotic Manipulators" and "Model Based Manipulator Control." Also see other publications: William Whittaker and coauthors' (1985) "Mine Mapping by a Robot with Acoustic Sensors," which Oppenheim also contributed as well as Robert Woodbury and Irving

Seminars at the Institute led Chris Hendrickson to make a similar shift.[59] In collaboration with his Ph.D. student Carl Haas, he focused on vision-based automation of crack sensing and sealing on highway slabs. An example outcome was a prototype vision system that Haas (1990) developed in the late 1980s.[60] The other notable example was William Whittaker and Esfandiar Bandari's (1986) concept system architecture for a multi-robot on-site connection workflow.

Whittaker and Bandari's (1986) system is important to reflect on, because it encapsulated the aspirations for construction built upon robots at the DCE in a similar way that the Flexible Manufacturing Cell/TRANSCELL encapsulated the ones for manufacturing at the CMU-RI. What we see in this technical proposition is a reformulation of on-site construction into a robot software operation, which "structur[es] the apparently unstructured" in Whittaker's (1986) words and makes on-site productivity more precisely manageable. The system includes a "host" that holds a global model of the real-world operation, and runs "global planning and optimization" of it by defining "sub-goals" based on the design, site, and "status of operations in progress" (Whittaker & Bandari, 1986) (Figure 10). The host, which Whittaker and Bandari (1986) calls as "counterpart to a human superintendent" is connected to a robot fleet they called as "subordinates." Based on the sub-goals received from the host, each robot performs obstacle detection, generates a local model, executes tasks on the site via navigation and manipulation, and sends the "dynamic changes" (position, progress, and changes in surroundings) to the host to update the global plan (Whittaker & Bandari, 1986). This makes the work conveniently trackable on the software. The system's human interfaces enable overseeing it, and interacting with the system when needed (Whittaker & Bandari, 1986). In this new configuration, the worker becomes the operator monitoring the workflow and productivity.

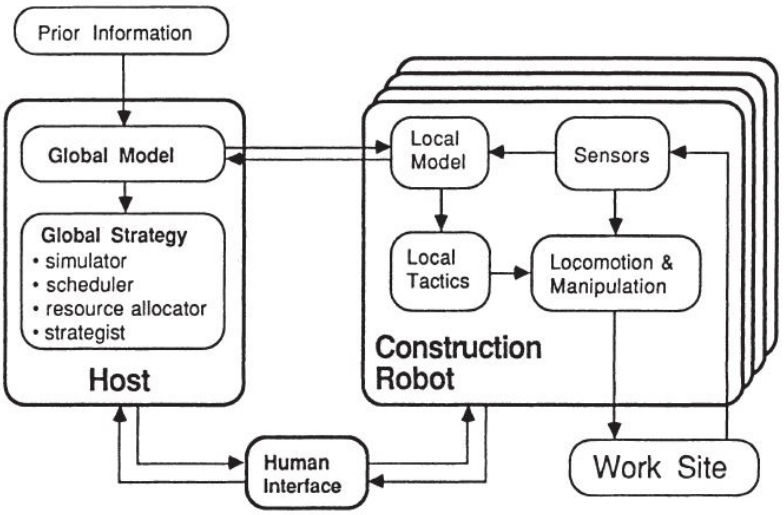


Figure 1: System Architecture

Figure 10. Whittaker and Bandari's (1986) system architecture for multi-robot construction (Source: CAD and Robotics in Architecture and Construction, Springer Nature BV)

Oppenheim's (1988) "An Approach to Geometric Reasoning in Robotics," and Darcy Bullock and Irving Oppenheim's (1992) "Object-oriented Programming in Robotics Research for Excavation."

59 Based on my interview with Hendrickson in 2021.

60 For details of this work, see Carl Haas and coauthors' (1992) conference proceeding titled "A Field Prototype of a Robotic Pavement Crack Sealing System" as well as Haas's (1990) PhD thesis titled "A Model of Pavement Surfaces."

Interactions with the Robotics Institute, the excitement for the possibility of robotization of construction raised by the two mobile robots, and the emerging research agenda intersecting robotics and on-site building tasks, on-site construction robotics emerged as an academic research field at the DCE in the 1980s. Despite these efforts pioneering a field, robot capabilities were yet to match the needs of tasks and sites, funding or interest from the industry was not really present, and no significant real-world deployment was in place in the United States during these early years.[61] Neither robots were fully ready for construction, nor construction was ready for robots. For that reason, the work at the Department was about situating the robots in construction technically, but more so through the discourse that I unpacked.

2.4.4 Decomposing Construction Work for Robot Development

While technical research was running at it natural pace at the DCE, a trajectory for robot adoption was being paved faster in the engineers' discourse than on job sites. There was a gap between the aspirations and realization within the formative years of construction robotics at the Department. The two studies by Warszawski and Sangrey (1985) and Skibniewski (1986) that I closely examine in this section broadly aimed at closing this gap, in other words, the gap between building tasks, sites, and systems existing in the early years of construction robotics. Motivated to address this gap, each study produced an unconventional technology design instrument: a taxonomy of construction work to identify robot performance requirements and a modular robot design framework, which are practical guides for robot development and adoption.

As I discuss in the following sections, I see these instruments demonstrating a second technology design approach shaped at the DCE, focusing on identifying, decomposing, and classifying construction tasks in service of robot development and adoption. This is distinct from the CECRL's approach seeking to advance robots' missing capabilities for construction. This second approach takes construction work as an object of analysis to inform robot development and adoption. It recognizes that converging construction and robotics was a truly interdisciplinary enterprise, in which the knowledge of practice and systems should effectively be bound together, and stakeholders in the industry should be informed. The role of the taxonomy and the framework that I see is to inform the two practices and their actors. Drawing from Geoffrey Bowker and Susan Leigh Star's (1999) definition, they are "artifacts that may … span highly complex boundaries," in this case the ones existing between construction and robotics practices in the 1980s.[62] Nonetheless, I point out that these two practices were still distant from each other at the DCE.

Similar to the high-level descriptions highlighting the engineers' recognition of the distinction between construction and manufacturing settings, I identify that this approach, along with the two unconventional technology design instruments that it produced, show the engineers' conception of construction work as a complex, layered makeup of activities as early as the 1980s. Both the taxonomy and the framework involve a high-level analysis of construction tasks, focusing on their general character as well as commonalities between them. They sought to find repetitive activities in each task, and similarly, repetitive activities common across tasks in order to match those with a single-task robot configuration. The ultimate benefit would be to decrease development and adoption cost of these technologies per task. However, I foreground that this high-level analysis also overlooked the situated nature and idiosyncrasies of tasks.

---

61 The lack of funding and deployments in construction robotics in the 1980s is stated in (1) Warszawski and Sangrey's (1985) "Robotics in Building Construction," previously referred, (2) Skibniewski and Russell's (1989) "Robotic Applications to Construction" previously highlighted, and by (3) Chris Hendrickson in his interview with me in 2021.

62 STS scholars Geoffrey Bowker and Susan Leigh Star's (1999) definition in quotes is from their book *Sorting Things Out: Classification and Its Consequences*.

A Taxonomy of Construction Work: Examining Warszawski and Sangrey's Approach

In 1985, Abraham Warszawski and Dwight Sangrey (1985) published a study titled "Robotics in Building Construction." Motivated to combine robotics and building construction, the engineers through this study argued that the robotic technologies were close to matching the demands of tasks and sites, and the tasks were good fits for the deployment of these technologies. In doing so, they demonstrated construction work, when systematically approached and analytically decomposed into its parts, appears to be well suited to robotization. They discussed that each part can be performed by at least one or more types of robot configuration (Warszawski & Sangrey, 1985). I see this effort as a way to promote the possibility of robots to the construction industry that was often regarded inefficient and slow in technology adoption. I also consider the engineers' demonstration of the match between robots and construction tasks a response to the lack of interest in construction robotics that they noted in the article. As the engineers (1985) highlighted, there was nearly twenty million in government funding for robotics in 1982 next to "almost no funds" in construction robotics.

Warszawski and Sangrey engaging both ends of the field shows the engineers' acknowledgement of construction robotics as a field that is highly domain dependent. The reason for the "lack of active interest" was construction being a complex and, hence, a challenging application domain with "unique nature in every project," "production moving from one location to another," and "rugged environments" as the engineers (1985) pointed out. Without an understanding of both work and technologies, raising interest or realizing the vision of putting robots into construction would virtually be impossible. Towards that understanding, the engineers (1985) combined a state-of-the-art review of system components and a two-step analysis of construction tasks.

Through the technical review, Warszawski and Sangrey (1985) assessed the applicability of robots for construction work.[63] The review involves identifying the technical requirements for each component to suit indoor and outdoor construction tasks, and evaluating whether the state of the art met those requirements. For example, for arm-based manipulators for handling indoor construction tasks, the engineers specified a ten-to-thirty-kilogram payload based on an average worker's "load lifting capacity," along with a three-to-four-meter reach, that is, "larger than that of an average worker." The engineers deemed that the industrial arms' payload at the time met the requirement, but their reach needed further improvement. Regarding end effectors, they concluded that a good variety of grippers for pick-and-place as well as specific tools, such as welding guns, sprayers, and drillers, were available for integration. Considering site conditions and safety, the engineers suggested "preprogramming" as the appropriate method of control and "well-structured portions of large tasks" as suitable for preprogramming. They identified "tactile," "proximity," and "vision" sensors as the requirements for operating in unstructured construction contexts, but without any technical specifications. Lastly, they listed teleoperation, preprogramming, and autonomous motion, which was quite primitive at the time, for robot navigation in construction.

Through the analysis of construction tasks, Warszawski and Sangrey (1985) produced a taxonomy-like classification of construction tasks, and based on the classification, discussed the suitability of tasks for robotization (Figure 11).[64] In this classification, the engineers first categorized the labor involved in different building components as tasks. Framing, formwork, attaching hangers, or

63 This review relies on various other sources, such as Joseph Engelberger's (1980) *Robotics in Practice*, Edward Fisher and Shimon Nof's (1982) "The Work Abilities of Industrial Robots," and Yukio Hasegawa's (1983) "Robotization of Construction Work" among others. The description of the technical review in this paragraph is based on Warszawski and Sangrey's (1985) article.

64 This analysis Warszawski and Sangrey (1985) described in "Robotics in Building Construction" was originally conducted and published by Abraham Warszawski (1984) as part of the first ISARC symposium (titled "Application of Robotics to Building Construction"). The description of the analysis of tasks in this paragraph is based on Warszawski and Sangrey's (1985) article.

painting are example items. Based on this categorization, the engineers then broke down these tasks into ten types of activities, such as positioning, connecting, or coating, each of which may be handled by a single robot configuration. Next, they examined these ten activities in terms of the manipulator, end effector, control, sensing, and mobility requirements needed in a system to perform them. Based on the system requirements for each activity, The engineers concluded that all ten activities—and, consequently, the tasks they are associated with—could be handled by only four types of archetypal robot configurations, characterized in the article as the “assembly (assembling) robot,” the “interior general-purpose robot,” the “floor finishing robot,” and the “exterior wall finishing robot” (Figure 12).

TABLE 1.—Basic Activities in Building Construction

No. (1)	Activity (2)	Description (3)	Examples of application (4)
1	Positioning	Placing a large object at a given location and orientation	Erection of steel beams, precast elements, formwork, scaffolding
2	Connecting	Connecting of a component to an existing structure	Bolting, nailing, welding taping
3	Attaching	Positioning and attaching of a small object to an existing structure	Attaching hangers, inserts, partition boards, siding sheathing
4	Finishing	Applying continuous mechanical treatment to a given surface	Trowelling, grinding, brushing, smoothing
5	Coating	Discharging a liquid or semiliquid substance on a given surface	Painting, plastering, spreading mortar or glue
6	Concreting	Casting of concrete into molds	Casting of columns, walls, beams, slabs
7	Building	Placing blocks one next to or on top of the other with a desired pattern	Blocks, bricks or stones masonry
8	Inlaying	Placing small flat pieces one next to the other to attain a continuous surface	Tiling, wood planks flooring
9	Covering	Unrolling sheets of material over a given surface	Vinyl or carpet flooring, roof insulation, wall fabric
10	Jointing	Sealing joints between vertical elements	Jointing between precast elements, between partition boards

Figure 11. Warszawski and Sangrey’s (1985) analysis of construction work (Source: The American Society of Civil Engineers)

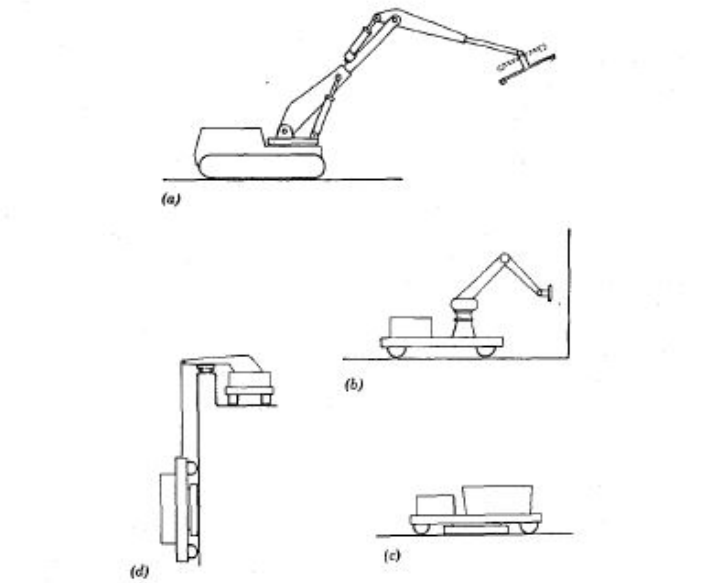


FIG. 3.—Construction Robots: (*a*) Assembly Robot; (*b*) General Purpose Robot; (*c*) Floor Finishing Robot; (*d*) Exterior Wall Robot

Figure 12. The four robot configurations (Warszawski & Sangrey, 1985) (Source: The American Society of Civil Engineers)

I see Warszawski and Sangrey's (1985) taxonomy as more than just a technical guide, foregrounding the match between tasks and robot configurations. Borrowing Bowker and Star's (1999) definition, it is also an "object[s] for cooperation across social worlds," in this case, providing a common framework, a standard identification of tasks through classification and common categories in relation to system components. The aim was to inform technology designers, researchers, and stakeholders.[65] Warszawski and Sangrey (1985) knew the importance of stakeholders' views for robot adoption, as they wrote robots should "win the support not only of labor but also of management at all levels." In that sense, the taxonomy was an unconventional technology design instrument geared to draw the different social and knowledge worlds of these actors closer. However, the worlds of robotics and construction were still seen separate than intertwined by the engineers. The technologies conceived and produced in robotics were considered as unavoidable disruption in job sites. This was displayed by the way the engineers (1985) foresaw "a considerable reorganization of the building work," "displacement of some labor," and "redefining the task of others" as default consequences of robotization in a later section of the publication.

Further, the engineers' decomposition of tasks into parts and classification of parts that repeat across tasks show their conception of work as a complex constitution of activities that needs to be analytically approached and broken down. It shows the engineers' focus on tasks' general character and commonalities. Idiosyncrasies of tasks, more specifically, variations in the way the common activities would possibly get enacted in different tasks and contexts are deliberately abstracted out.

A Modular Robot Design Framework: Examining Skibniewski's Approach

A year after Warszawski and Sangrey's (1985) study, Skibniewski (1986) published his dissertation, titled "Engineering and Economic Analysis of Robotics Application Potential in Selected Construction Operations." The dissertation offered a cost-benefit evaluation framework for robot applications in construction and, based on this framework, reported on the potential cost-benefit impact of robotic concrete sandblasting and concrete formwork cleaning applications with numbers (Skibniewski, 1986). Under each application case were the "robot design specifications" that Skibniewski (1986) outlined as a basis for his evaluation. This part of his project was similar to Warszawski and Sangrey's demonstration of the match between tasks and robot configurations but more detailed in its descriptions of system components, workflow, and a conceptual system for each of the two selected tasks (Figure 13). Ultimately, both works promoted the possibility of robotization in the industry, but in slightly different ways. While Warszawski and Sangrey's work did so by arguing that the existing technology could in theory address all tasks, in Skibniewski's evaluation, it was by arguing for the value that the two conceptual systems could bring into sandblasting and cleaning, such as saving costs, increasing productivity, and improving quality.

Although the two cost-benefit evaluation make up the bulk of Skibniewski's (1986) thesis, a short section titled "Prospects for a Modular Design of Construction Robots" proposes a distinct technology design concept for the young and growing field of construction robotics. In this section, Skibniewski (1986; Skibniewski et al., 1986) described a modular approach to robot design built on the premise that a machine constructed with modular components could perform a group of similar tasks. According to him, "similarities in design approaches" to single-task systems could be studied and streamlined towards building one modular multi-purpose machine, which would provide value in more than one task (Skibniewski et al., 1986). With that, the engineer attempted to show construction work, when analytically decomposed into its parts, can help capture similarities between

65 Bowker and Star's (1999) definition of "classifications" as "objects for cooperation across social worlds" is included in the introduction chapter of their book *Sorting Things Out: Classification and Its Consequences*.

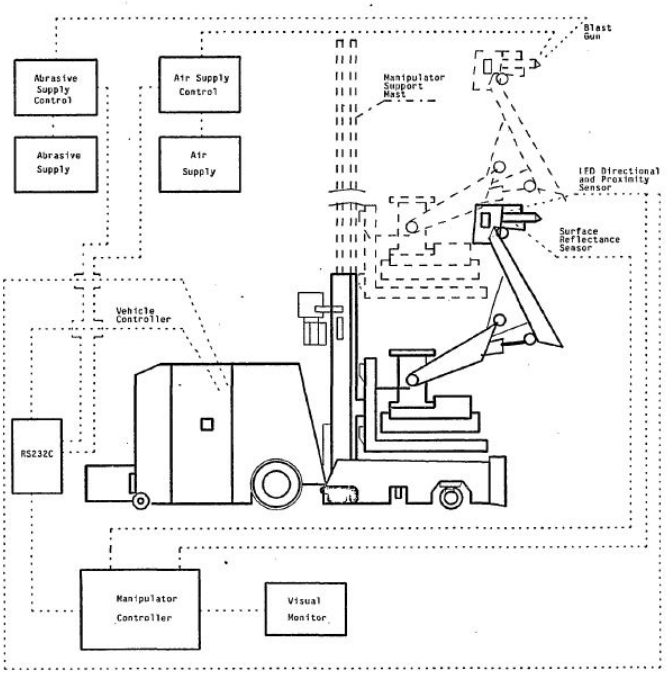


Figure 3-2: General Setup of the Sandblasting System.

Figure 13. Skibniewski's (1986) conceptual system for concrete sandblasting (Source: Miroslaw Skibniewski)

tasks, and ease the cost barrier of robotization by distributing it over those similar tasks. To help readers imagine this, Skibniewski (1986; Skibniewski et al., 1986) used the example of surface tasks, which, he suggested, "can be performed with similar robotic control strategies"; nonetheless, he did not provide an example of what a modular system for these tasks would look like. Instead, the engineer (1986) introduced a two-part design framework for governing the making of such modular systems.

The first part of the framework is a "robot purpose analysis," which can be summarized as a systematic analysis of construction work and site requirements towards producing robot design implications (Skibniewski, 1986). (Figure 14).[66] The analysis starts with the "definition of robot work functions," which lays out each group of construction tasks, such as "structural," "surface," or "maintenance and cleaning" operations. It is followed by the "identification of individual work tasks," essentially a list of tasks in each group, classified according to their common characteristics. The next step of a robot purpose analysis is the "identification of various work site locations," that is, categorizing the locations on a job site, such as the "interior and exterior of buildings," and extracting the requirements associated with each location. The goal is to reach a set of system components satisfying as many location requirements as possible. Because the selection of these components is also contingent on "environmental and performance parameters," the next step consists in the "establishment" of these "performance parameters," such as the reliability of components with respect to "environmental conditions." The final step sees the "identification of field hardening requirements" to further increase the system's reliability.

66 The two parts of the robot design framework, each of which includes "several stages of activities" as Skibniewski (1986) describes, are based on Horst and coauthors' (1984) "Procedure for Developing a Standardized Robotic Maintenance System for Nuclear Power Plants" cited in his dissertation. The description of the first part of the framework in this paragraph is based on the dissertation.

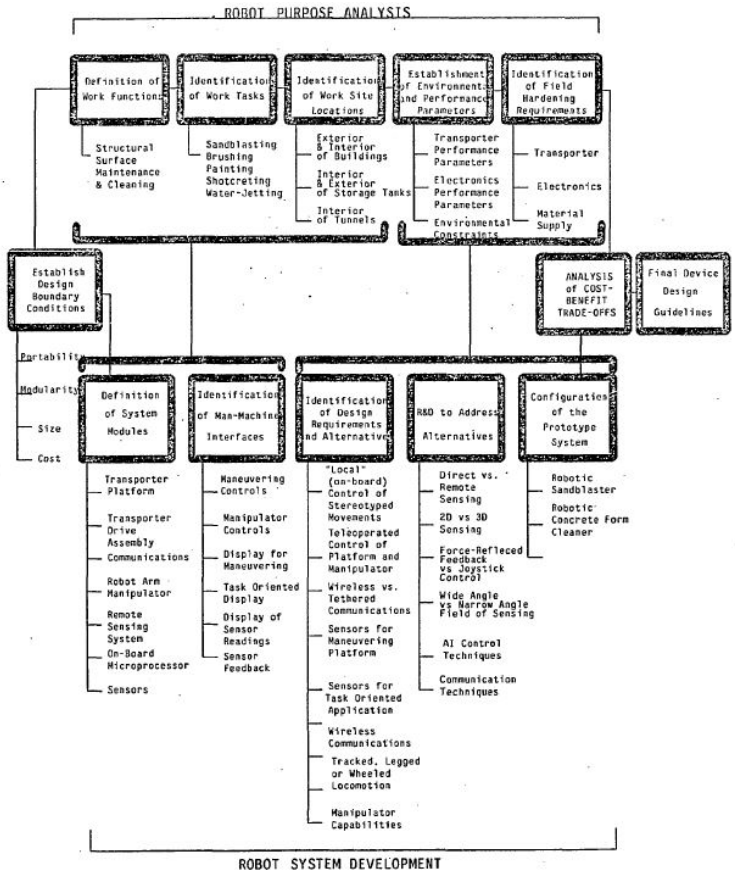


Figure 2-1: A Modular Approach to Robotic System Design.

Figure 14. Skibniewski's (1986) two-part modular robot design framework: (1) robot purpose analysis and (2) robot system development (Source: Miroslaw Skibniewski)

Based on the robot purpose analysis, the "robot system development" follows as the second part of Skibniewski's (1986) framework (Figure 14).[67] The first step in this development is the "definition of the system modules," such as mobile base, manipulator, or sensors, that fit the requirements of the construction tasks identified and grouped during the purpose analysis. The second step focuses on the "identification of man–machine interfaces" in the system, such as controllers, displays, or sensor feedback, which could maximize flexibility for essential human intervention. System modules and interfaces to implement are all identified through careful consideration of the task and work site location requirements. The next step in the development, the "identification of design requirements and alternatives," concerns evaluating the "efficiency," "feasibility," and "availability of [technology] solutions" for implementation, such as "tracked, legged, or wheeled locomotion" or "manipulator capabilities." The "research and development necessary to address the alternatives" comes next for the further improvement of system capabilities as new solutions in sensing and control emerge. The last step consists of the "configuration of the prototype system" guided by all these steps.

Similar to Warszawski and Sangrey's (1985) taxonomy, Skibniewski's (1986) modular design framework is also an "object[s] for cooperation across social worlds" (Bowker & Star, 1999), informing technology designers, researchers, and stakeholders that single-task robots could be

67 The description of the second part of Skibniewski's (1986) framework in this paragraph is based on the dissertation.

made economically viable. The framework is an early proposition for modular versions of these technologies that can handle multiple similar tasks, which meant an expansion in their function and a reduction in their costs.[68] While Warszawski and Sangrey postulated that almost all building tasks could be handled by four types of robot configurations, Skibniewski suggested that one modular configuration could be turned into several systems to handle multiple tasks. This would allow one investment for a machine distributed over multiple "conceptually similar applications," which Skibniewski (1986) presented as an incentive for the industry. Also similar to the taxonomy, the framework aimed at drawing the worlds of robotics and construction closer through this incentive by design, however, the abridging that it offered was only managerial.[69] Similar to Warszawski and Sangrey's work, Skibniewski's work reflected worker displacement as an inevitable consequence of robotic technologies. It featured displacement as an element in cost-benefit estimation potentially bringing cost-saving and productivity benefits to contractors (Skibniewski, 1986). For instance, a potential benefit of a robotic (concrete formwork) cleaning application would be cost-saving by "elimination of manual low-skill labor required to perform [the task]" (Skibniewski, 1986).

Lastly, like Warszawski and Sangrey's (1985) taxonomy, Skibniewski's (1986) modular design framework displays a conception of construction work as a complex constitution of activities. The analytical decomposition of this constitution, in Warszawski and Sangrey's case, resulted in a classification of tasks into ten activities that the state-of-the-art robot configurations could potentially handle. In Skibniewski's case, the decomposition of work produced a classification by common character, and this instrument was used to argue that robots on job sites could be an economically viable proposition.

2.5 Summary

Computer control and high-level languages were the key advances that facilitated a more practical reprogramming of robots from the 1960s to the late 1980s. These advances were adequate for arm-based manipulators to be deployed in the wild, as we see in the examples of PUMA and T3 in the late 1970s auto industry. In the early 1980s when the United States' position in international trade was under question, the arms' growing potential based on these advances created hopes for accelerating the production of diverse export goods at a lower cost. As I traced in the chapter, these hopes, combined with the Westinghouse partnership, cultivated a robot-productivity discourse central to the research at the CMU-RI. Boosting US batch production productivity with programmable automation and, consequently, the country's position in the global economic race were the elements of this discourse I emphasized in my analysis of the Ayres and Miller's (1981) report. Systems combining arms and high-level languages, such as the Flexible Manufacturing Cell/TRANSCELL developed by the Institute in collaboration with Westinghouse, were motivated by this discourse.

As I covered in the chapter, the mobile robots by the CECRL were the early examples that showed robots combining manipulation and navigation in construction could be a plausible idea. Alongside the excitement the two prototypes created, a group of civil engineers' knowledge exchange with the Institute and aspirations of building an agenda on robotization of on-site building tasks led to the emergence of construction robotics as a research field at the DCE. While robotics at the CMU-RI concerned the transition from hard to programmable automation with robots in the manufacturing

68 Even though Skibniewski (1986) did not give any clear description of how system modules get assembled or reconfigured to switch between similar tasks in the dissertation.

69 For instance, the data that allowed Skibniewski (1986) to estimate the costs of sandblasting or cleaning applications was primarily drawn from building contractors in Japan with prior system development experience, contractors in the United States and Germany with domain knowledge of these tasks, and companies outside construction in the United States and France with systems applicable to these tasks.

industry, these engineers, aspired to realize a similar robot-work integration in construction, were situating these technologies into this new work domain both technically and discursively. Their agenda comprised not only technical research, but perhaps built with a higher pace than technology, aspirations to reconfigure work in the construction industry. As I revealed in the chapter, these aspirations were built around a robot-productivity discourse similar to the one at the Robotics Institute as well as high-level descriptions of construction domain, foregrounding the differences between job sites and manufacturing settings to emphasize the challenges in construction for robots, but also the similarities between the two work domains to make a case for these technologies.

Based on these two elements and systems described, I pointed out that the civil engineers built excitement around robots, the same technology that had shown promise in manufacturing automation and in the CMU-RI projects, and a discourse around robots that similarly underscored productivity. However, unlike manufacturing, the domain that they situated these had no previous layer of automation, which meant a bigger transition from hand tools to robots. I also identified a recognition of construction as a complex, layered makeup of site conditions indicated by the construction-manufacturing distinctions they made, and a conceptualization of certain tasks as more similarly repetitive to manufacturing tasks than they are, shown in some of their work. A finer grain repetition arranged by the existing machines and contrived settings in manufacturing compared to broader repeating cycles of a larger mix of actions and numerous improvisations on job sites might be overlooked. As I also showed in the chapter, the material-oriented character of a set of tasks and workflows (critical to account for in developing technologies) were not part of the understandings of work as the focus was on autonomy, sensing, or motion needed to perform those with machines.

In the chapter, I also zoomed in on Warszawski and Sangrey's (1985) taxonomy of construction work to identify robot performance requirements, and Skibniewski's (1986) modular robot design framework. I identified a recognition of construction work as a complex, layered makeup of activities in these engineers' approach of identifying, decomposing, and classifying building tasks through these instruments. I highlighted that the engineers were interested in the general character and commonalities of building tasks. The aim was to match each or a set of tasks with a robot configuration in a cost-efficient manner. However, their high-level analyses seeking to find commonalities overlooked the situated nature and idiosyncrasies of tasks as I discussed.

Further, I showed that the work at the DCE displayed two technology design approaches in the emergent years of construction robotics. The first one that I discussed through the CECRL's prototypes focuses on developing technologies to fulfill the missing affordances that machines needed to work in construction. The motivation engraved in this forward-looking approach is to demonstrate the possibility of robotization of construction through showcasing system capabilities. The second approach I identified in Warszawski, Sangrey, and Skibniewski's work focuses on evaluating if construction work and technologies are well-suited for each other. Through this approach, analyzing work appears as an alternative way of showing the possibility of robotization next to showcasing system capabilities. I situated this approach as an attempt to pull the two knowledge worlds of robotics and construction.

The vision, efforts, and work I described in this chapter positioned the DCE at the driver's seat of construction robotics in the United States. By the end of the 1980s, most of the construction robot prototypes available in the country were developed at the Department and the Robotics Institute (Skibniewski & Russell, 1989) (Figure 15). However, these systems mostly stayed at the lab, and had attracted neither any significant resources nor the construction industry's interest for further development or deployment. The Department pushed for the industry's first foray into automation,

yet the existing infrastructure was not yet ready for robotic technologies. Thus, as construction robotics was emerging as an institutionalized field at 1980s CMU, how far construction remained from the vision of automation was a question buried underneath the excitement around robots. The aspirations were much more ambitious than the construction sites and robots could handle. Building robots to achieve the missing capabilities and analyzing work to close the gap between robots and the reality of construction were running parallel at the DCE. However, these two approaches were still divergent, and had not yet cohered into practice-based technology design practice.

TABLE 1 — EXAMPLE CONSTRUCTION ROBOTIC PROTOTYPES DEVELOPED IN THE UNITED STATES

System Description (1)	Application (2)	Research Center (3)
John Deere 690C	tele-operated excavation machine	John Deere, Inc, Moline, IL
Laser-Aided Grading System	automatic grading control for earthwork	Gradeway Const Co & Agtex Development Co, San Francisco, CA; Spectra-Physics, Dayton, OH
Automatic Slipform Machines	placement of concrete sidewalks, curbs, and gutters	Miller Formless Systems Co, McHenry, IL; Gomaco, Ida Grove, IA
Micro-Tunnelling Machine	tele-operated micro-tunneling	American Augers, Wooster, OH
Robotic Excavator (REX)	autonomous excavation, sandblasting, spray washing & wall finishing	The Robotics Institute, Carnegie Mellon Univ, (CMU), Pittsburgh, PA
Autonomous Pipe Mapping	mapping subsurface pipes	The Robotics Institute, Carnegie Mellon Univ, Pittsburgh, PA
Terregator	autonomous navigation	The Robotics Institute, Carnegie Mellon Univ, Pittsburgh, PA
Remote Core Sampler (RCS)	concrete core sampling for radiated settings	The Robotics Institute, Carnegie Mellon Univ, Pittsburgh, PA
Remote Work Vehicle (RWV)	nuclear accident recovery work, wash contaminated surfaces, remove sediments, demolish radiation sources, apply surface treatment, package & transport materials	The Robotics Institute, Carnegie Mellon Univ, Pittsburgh, PA
Wallbots	construction of interior partitions metal track studs	Massachusetts Institute of Technology (MIT), Cambridge, MA
Blockbots	construction of concrete masonry walls	Massachusetts Institute of Technology, Cambridge, MA
Shear Stud Welder	weld shear connectors in composite steel/concrete construction	Massachusetts Institute of Technology, Cambridge, MA
Automated Pipe Construction	pipe bending, pipe manipulation & pipe welding	University of Texas at Austin, Austin, TX

Figure 15. Construction robots developed in the United States by the end of the 1980s (Skibniewski & Russell, 1989) (Source: American Association of Cost Engineers)

Chapter 3

Studying Construction Work Towards Robot Development: The ReAC Project

3.1 Outline

Between 2021 and 2023, I actively participated in a research project by CMU's ReAC group, in which we explored a practice-based approach to designing a robotic system for construction. With the ambition of acquiring a richer understanding of construction work, the project incorporated a fifteen-week ethnographic study into a technical exploration of RL techniques for construction robotics. I designed and conducted this ethnographic study at the construction site of a new educational building called the New Scaife Hall on the CMU campus. This study aimed at bringing the knowledge of tasks and workers lying on the site into conceiving and developing a system. To trace how this knowledge transaction plays out, I conducted a second ethnography of the group's system development activities, which mainly took place at a CMU lab.

In this chapter, I present the produce of the first ethnography, which I call the elements of the study, namely, the empirically rich descriptions of building tasks, the knowledge of zones and interactions, and the workers' perspectives on robots. In the next chapter, I focus on the second ethnographic study, and trace how ethnography drew robotics and construction closer by identifying the ways that these elements guided the group's conception and development of a system. Further, I trace our conceptions of construction work and automation as they were shaped by these elements.

In the rest of the chapter, I unpack the elements of the New Scaife Hall study. I first present the task descriptions to show the depth and richness of building tasks, which we were interested in understanding and engaging with in system development. Then I share the knowledge of zones and interactions, which helped us to comprehend how work is conducted on a construction floor through a list of points on organization, movement, and interactions. First, materials were spatially organized in designated zones based on the building trades' task location and coordination, and restrained movement in construction. Similarly, trades performed tasks in virtually bound zones based on overall planning. Workers' movements were mostly confined within these zones to maximize safety and minimize conflicts between each other. Lastly, interactions on the floor were more common within a unit or a trade, whereas they were limited to coordination purposes across them.

Next, I offer a synthesis identifying the qualities of the observed tasks and offering key insights by which the group viewed them. The first insight is that certain qualities are only visible on a construction site or when conversing with workers. The second one is that tasks or activities possess general qualities to different extents. Further, tasks entail different degrees and paces of mobility along with different degrees of collaboration and communication, dictated by the role of workers and materials. Finally, I discuss the workers' perspectives on robots. The workers have mixed views on the ambition of robot adoption in construction in the United States. Their views, which we carefully considered in conceiving a system, concerned the situated nature of tasks, skills and deskilling, fitness of robots and sites, and the appropriate roles for robots in construction.

3.2 An Ethnographic Study of Construction for Robot Development

3.2.1 The New Scaife Hall Site

On a cold February morning, I was walking on the same Schenley Park grounds where the Terregator was demoed almost forty years before, about to reach my new workplace as part of the ReAC project. My destination was CMU's New Scaife Hall construction site, which I will also refer to as "the site" in this chapter. Located in the southwest corner of the campus, this building would be next door to Hamerschlag Hall standing gloriously as the University's landmark. The building under

construction would replace the original Scaife Hall built in 1962 with a brand new eighty-five-thousand-square-foot cross-disciplinary AI/robotics research facility (Gerson, 2020; Carnegie Mellon University College of Engineering, 2022). Initiated around the vision of "propel[ling] forward new research fields, such as intelligent systems, Softbotics, and drone research," this new building was on the way to becoming a contemporary emblem of CMU's tech-entrepreneurial ambitions around robotics, which also drove the foundation of the Robotics Institute, as I described in the previous chapter (Carnegie Mellon University College of Engineering, 2022). Before serving this vision as a complete facility, the building under construction was about to contribute to robotics in a way that the University had not planned for.

For the next fifteen weeks, the New Scaife Hall construction site would serve as an ethnographic site where I, as an active participant in the ReAC project, would engage with construction work and workers to guide our design and implementation of a prototype robotic system. As I walked closer to the site, two large steel frames delineating the silhouette of the future facility came majestically into view. I stood at the site's entrance on Frew Street along Schenley Park, carefully stared at these frames, and made a quick mental map aided by the construction plan drawing in my hand (Figure 1). The site is located on a hill, with its lower grade facing the Pittsburgh railway, Schenley Bridge, and the Oakland neighborhood's modest skyline to the west, whereas it is at grade with Frew Street, with Hamerschlag Hall, ANSYS Hall, and Schenley Park surrounding it on the north, east, and south, respectively. At the lower grade, the facility's two-story podium sits beneath its towers, which is hidden when one approaches from the campus or the park but exposed to any viewpoint from the west. At the time, a series of trailers were lined up on the site's northern and eastern edges, in which the general contractor and subcontractor teams stored part of their tools, planned their work, and rested during their breaks.

The New Scaife Hall site was made accessible to our group by the general contractor and the University under a research collaboration. After I oriented myself towards the entrance, I walked to the northeast corner to meet with our site contact, a project manager from the contractor team. Shortly after, he arrived with a site engineer to get me in through the safety fence. My first day on the site began as soon as I stepped in. Our first stop was the Superintendent's office, or what I call "the control room of construction," located in a trailer on the northern edge. Inside this surprisingly cozy space were two work desks with a desktop computer and office supplies, a larger table with large construction drawings, and a set of chairs that seemed to be used for impromptu meetings between the Superintendent and subcontractors. Human-size boards covered with neon-colored sticky papers, which indicated subcontractor job schedules on a calendar grid, stood against the walls, aptly illustrating the nature of this gentleman's job of managing actors and moving parts (Figure 1). In a conversation weeks later, the Safety Manager would use the word "priest" as an analogy to describe the Superintendent's role on the site. I was formally introduced to the "priest" of Scaife Hall and about to enter his "church." The Superintendent reminded me about the site protocols before the site contact, the engineer, and I left the trailer for our next stop.

To minimize my footprint on the site, I chose to limit my observations to a single floor. The two podium floors, Level A and B, were the closest prospects because they would be the first floors in the project with a concrete deck in place for walking safely. At the trailer, the site contact informed me that Level B's deck would be completed very soon, so it sounded like this would be the floor for my observations. To show me the floor, he and the engineer escorted me to Level B. As I mentioned earlier, the podium floors sat on the lower grade, and a temporary entry for workers' access to those

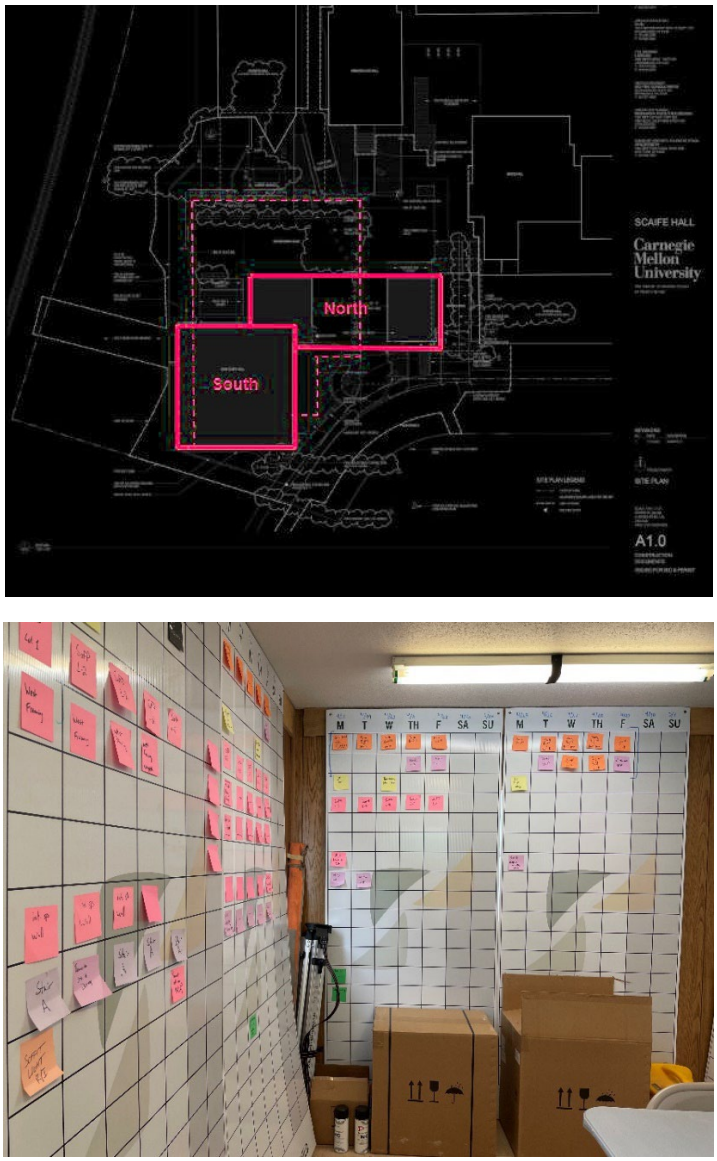


Figure 1. The New Scaife Hall plan (top) and scheduling boards inside the Superintendent's trailer (bottom)

floors was set up on the northern side of the site, slightly downhill from the Superintendent's trailer. From the trailer, we walked down a path covered with a mix of earth and gravel and crossed a temporary bridge and a temporary staircase, both built out of scaffolding, to reach Level A first. Continuing through another temporary staircase, we moved down to Level B. As we arrived, the tool and equipment sounds of ongoing concrete pouring, along with the high and at times nervous tones of workers pouring it, became vivid. The floor plan in my hand showed around nineteen thousand square feet (one thousand seven hundred forty square meters) of area in a fat L-shaped layout, circumscribed by a combination of structural walls and a retaining wall (Figure 2). At that stage of construction, the floor had no framing or interior partition and was open to the outside on its west side, with no framing or external wall (Figure 2). "Trades would start from the north end [of the floor] and move towards south," the site contact noted before he left with the engineer. I spent another hour on the floor testing the cameras and camera positions accordingly. All were ready for the start of the study the next day.

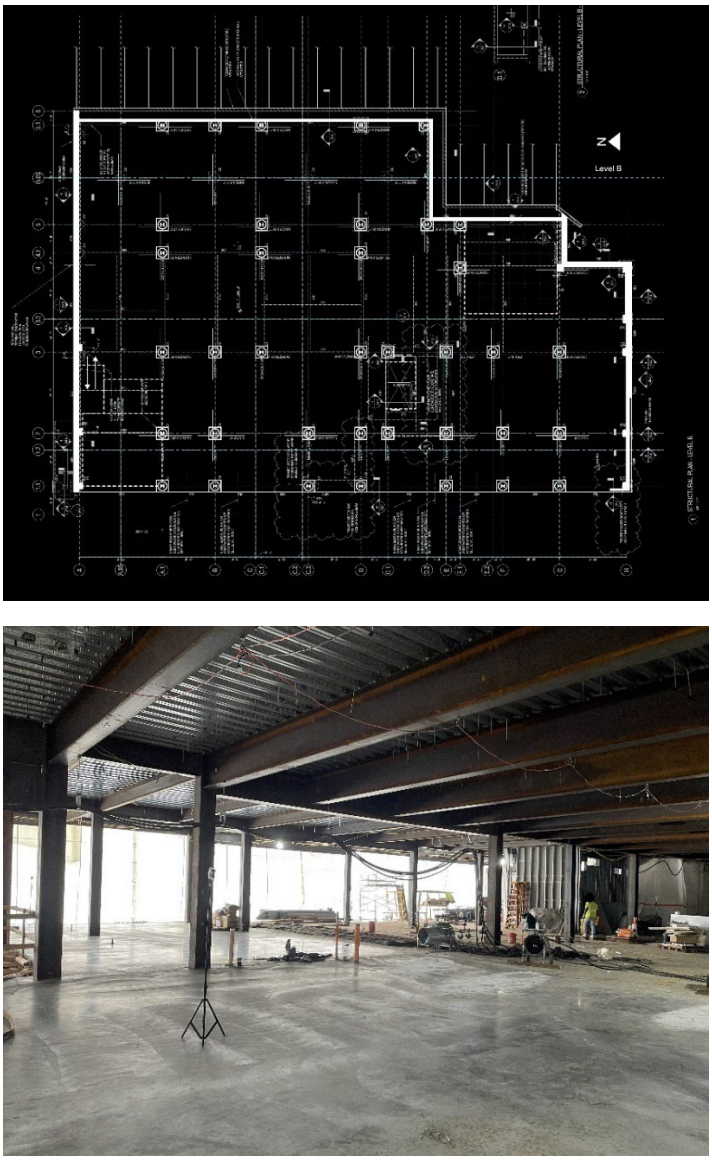


Figure 2. New Scaife Hall's Level B plan (top) and Level B on the first day of the study (bottom)

3.2.2 Descriptions of Building Tasks[1]

Through observations, time-lapse recordings, site documentation, and interviews, I collected a large ethnographic data set including a set of construction tasks, namely, cleaning-housekeeping, fireproofing, rebar tying, carpentry formwork, MEP ceiling line installations, framing, and drywall installation that took place on Level B. Based on weekly analysis of this data, I synthesized empirically rich descriptions of these tasks. In biweekly gatherings called site information sessions, I presented these descriptions to the group accompanied by the data viewer that I developed, displaying scenes and information from the site.

1 Because the interviews came after the observations, the initial versions of these descriptions did not include workers' points and descriptions. The final versions featured in this chapter are enriched by those as well as a more systematic reflection on tasks' qualities based on the analysis of interviews.

This section offers four of these descriptions (as vignettes from the site) to convey the depth and richness of building tasks that we aimed to understand.[2] This depth and richness in these descriptions along with the scenes of enactment on the data viewer placed immersion, imagination, and conceptual guides next to technical knowledge to think with, which I discuss in detail in the next chapter. Learning about tasks through these descriptions, together with the knowledge of zones and interactions, and worker perspectives, fostered viewing and analyzing tasks in a more granular level, and consequently situating a robot in a specific task with an alternative role. I provide details of this process in the next chapter as well.

Cleaning and Housekeeping

> *I am not OCD [obsessive-compulsive disorder], it's just, I work there eight to twelve hours a day, and it [cleaning-housekeeping] just makes it a lot easier for everybody else. And it's not even for me, I am doing it to help everybody else. (Cleaning-Housekeeping Worker 1)*

Cleaning and housekeeping entails removal of garbage and leftover materials along with organization of reusable materials, tools, and equipment in appropriate locations on a construction floor. The aim is to maintain cleanliness and order for workers to proceed swiftly and safely. For that reason, this task, which is often not even known to outsiders, is crucial in construction. Cleaning and housekeeping on Level B was carried out by a single worker, Cleaning-Housekeeping Worker 1, to prepare the floor for the upcoming fireproofing and MEP ceiling line installations.

The site conditions in which cleaning and housekeeping could take place are contingent on the task that precedes it. In this case, it was concrete pouring, so the floor's southern half was open with a freshly clean concrete deck, and its northern half was contrastingly messy with garbage, leftover materials, and assorted construction equipment left behind. "Well, it's a lot of guys' lunches, drinks, cigarettes, just wrapping that is around materials, boxes," Cleaning-Housekeeping Worker 1 explained to me later in an interview. "You're constantly moving [while] doing concrete, so, people hurrying up, grabbing a bite to eat, throwing garbage on the ground, grabbing a drink, throwing it on the ground," he continued, not to blame the concrete crew he was part of but to point out the negligence and mess caused by the fast pace in construction. That day, I observed him cleaning and tidying up the messy northern half (Figure 3).

After spending hours with Worker 1 on Level B, it was obvious to me that his task relies on physical strength rather than craftsmanship and dexterity. In a nutshell, this worker spent those hours repetitively bending down, hand-picking, lifting, carrying, scraping, and shoveling items. His flow started with the western portion of the messy half. This is where I observed the cleaning part of his task, which consisted in removing the garbage and leftover materials from the floor and sorting the reusable ones. He began by clearing a bundle of leftover wood pieces, nylon sheets, and material bags he had previously piled on a pallet, a simple wood base to keep materials slightly above the floor. He reached under the pallet with the pallet jack, a wheeled cart with a steer to lift and move heavy items, and lifted the bundle.

Among Worker 1's activities, walking probably deserves a separate highlight because it characterizes cleaning as a highly mobile task. A big part of it is traversing between different areas and points he designated as garbage collection or material storage spots, at times with a load on the

2 The rest that describe MEP ceiling line installations, framing, and drywall installation are included in Appendix C.

pallet jack. Starting the first of his seven of those loaded walks that day, Worker 1 began pulling the jack to a collection spot he had set up in the clean half, which would later be a drone test arena in the finished building but was then a skylit courtyard where materials were vertically transferred down to Level B (Figure 3-A). He pulled the jack for about one hundred feet until he reached the spot, and, on the way, he bent down to hand pick or used his legs to move small obstacles out of the way. Once he arrived at the spot, he dropped the pallet loaded with the bundle on the floor and brought the pallet jack back. Following this round, he bent down and began working on a second pile, which was a mix of reusable wood pieces, some leftover materials, and litter (Figure 3-B). While he was carefully sorting the reusable wood, he collected the litter in a garbage bag. After sorting, he carried the wood to a spot he designated as its storage area. Lastly in this zone, he collected the litter—mostly plastic bottles and unusable hardware—into a bucket and dumped it into the elevator core, which he had used as a temporary dumpster at that stage of construction.

A salient quality of Worker 1's movements was that they were diverse, some repetitive over a period of time but unscripted. They were based on a mix of planning and improvisations, which is the intellectual layer in this task that is seemingly just physical on paper. Part of Worker 1's planning comes from the Superintendent: "I meet with my boss every day, he explains [to] me what's gonna be going on." This explanation includes the trades, materials, equipment leaving and coming as well as their transitions on the floor. For instance, the central and eastern portions, where I observed Worker 1's cleaning and housekeeping activities were where the fireproofing would start the next day, and they would soon get busier with the other trades. Thus, they should be relatively open and easy to move around. The eastern portion was the most congested with items used in concrete pouring, including a material bundle laying on a large cloth on the floor and two ride-on trowels (vehicles for smoothening concrete) parked in line. To tidy up, he first picked up some materials and gathered them on a pallet that was already set up (Figure 3-C). Then, using a hand truck, he lifted the first trowel off the ground, attached rubber wheels, and pulled the vehicle to the drone test arena, from where it would be transferred to the ground level (Figure 3-D). He walked back and repeated the same steps to move the second trowel. This was followed by moving a power trowel (the handheld version of the vehicle for smoothening concrete) and a heater to the drone test zone in two more rounds of pallet jack transfer. Each of these walks was approximately fifty to seventy feet long.

The other part of planning happens in Worker 1's mind on the floor: "I try to visualize how I want it to be," he said, and this visualization is a loose plan based on where the mess is, where he would collect and send the garbage, and where he would put the reusables for the trades coming in. For instance, the central portion included a bundle of rebars, which would be reused by ironworkers in a rebar tying job on the eastern edge. Later, in our interview, Worker 1 explained: "I kind of know who it [the material] goes to, so I find out where they are at, and I'll try to get it in that area."

How he would enact the plan is improvisational based on the floor layout and the type of mess. The use of pallets and pallet jacks as well as using the drone test arena and elevator core as garbage collection spots were some of Worker 1's creative improvisational solutions. With the rebars, he chose to use his hands. He would bend down, holding a bunch of rebars from their center of gravity, and lift them to waist level to transfer them to the eastern portion. Alternatively, he would drag them. With similar body movements, he gathered the wood beams and positioned them next to the cloth. When all that was left on the cloth was material waste, he folded the cloth to bag it and went back to using the pallet jack. He slid two long wood beams and a panel underneath this bag of waste and used the jack to pull it to the garbage spot (Figure 3-E). For the next series of activities, he switched to using a mix of tools and his hands. He was clearing the debris that accumulated in the eastern portion, which he combined with collecting the plastic bottles lying about in the central portion. He



Figure 3. Key moments of Cleaning-Housekeeping Worker 1's workflow on Level B

swept and gathered the debris with swift arm movements (Figure 3-F). With a quick transition to the central portion, he collected the bottles in a bucket and placed the bucket next to the garbage bag to be removed later. He then returned and shoveled the debris into two cardboard boxes and placed them on a pallet, also to be transferred from Level B later.

Worker 1 finished the day with some housekeeping back in the central and western portions. He sorted another bundle of rebars and wood beams and stored the beams on a designated spot in the western portion of the floor, followed by rebars and the remaining beams, which he added to the ironworkers' stack in the eastern portion (Figure 3-G). At the end of the day, all reusable items were compiled on the western half of the floor, while the eastern half was open for fireproofing and MEP ceiling line installations.

Fireproofing

> *The more that it [fire resistant] flows through the hose, the less problems you have… If you take too long of a break, if you have a problem with the pump or something, that stuff will start to set up in the hose… That's why it's so important to keep it going. (Fireproofing Worker 3)*

Fireproofing involves coating columns and beams of a building with a specialized fireproofing substance called the "fire resistant" based on the fire-rating classifications of these components. As Fireproofing Worker 3 who is the foreman of the fireproofing crew on Level B explained, these classifications refer to the number of hours that the component can withstand fire, and "the thicknesses [of the substance on columns and beams] are achieved to create that." The division of labor is between three workers: one lead worker on a scissor lift or on the ground spraying the fire resistant with the "spray gun," a second worker running the mixer that prepares and pumps the resistant to the gun, and a third worker with a supporting role performing concurrent cleaning, switching on and off the flow between the mixer and the gun, and managing the hose (connecting the two devices) that may obstruct the traffic on the floor.[3] The portion of the task that I observed was conducted by Fireproofing Workers 1, 2, and 4 within a perfect rectangle demarcated by four column lines on the northern half of Level B (Figure 4).

Fireproofing is infamously messy and requires gradual movement between columns and beams, and therefore demands a relatively more open zone distant from other tasks. As Level B entered my field of vision, a large puddle with a spread of gray, cementitious fire resistant debris on a sizable area in the northwest displayed the mess the task typically generates. Two thick hoses curled next to the puddle extended to a scissor lift standing close to the debris. On the lift's platform was Fireproofing Worker 1 in a white jumpsuit, spraying a steel beam with the spray gun, which is a long, stick-like steel dispenser with a trigger connected to the end of the hose. At the other end of this seven-hundred-foot hose was a mixer at the ground level, run by Fireproofing Worker 4. A large portion of the resistant from Fireproofing Worker 1's gun stuck to the steel surface while part of it sprinkled in a fireworks pattern adding to the debris. Away from the lift, Fireproofing Worker 2 at the supporting role was sweeping the areas where he had recently sprayed with a squeegee.

For anyone watching Fireproofing Workers 1 and 2, the unceasing continuity of their movements is the quality that catches the eye. This is not only intended for efficiency but also conditioned by the fire resistant. Once the substance enters the hose, it must leave soon lest it starts drying. Besides,

3 Based on my interview with Fireproofing Worker 3, after being prepared in the mixer, fire resistant is pumped through the hose to a particular floor and finally injected with air, which turns it into a spray.

once the substance falls on the floor, it should be cleaned before it fully dries. “It is ten times harder to clean it [the resistant] up [when it dries],” Fireproofing Worker 3 noted. Therefore these workers have to “keep it going,” as Fireproofing Worker 3 explained. In a way, the two workers have to adapt their pace to a machine’s flow, or, in exact terms, the fire resistant’s drying speed in the hose or on the ground. Doing this for long periods demands efficiency, which I noticed the workers established through the mechanics of their respective routine movements over time. Pacing their routines with the flow also demands physical strength and endurance needed to sustain them.

With Fireproofing Worker 1, it was his left-right routine with his arms slightly raised above his shoulders that enabled him to cover lanes with few passes. The routine also evidenced his dexterity as he controlled his arms’ pace and distance relative to the viscosity and amount of fire resistant that stuck to steel. When I observed this, he was on the lift perpendicular to a beam running north to south and delimiting the rectangle’s western edge. His head was close to the ceiling, and the beam’s lowest part was around his chest. Facing the beam, he sprayed its middle portion with a subtle, left-and-right movement with his elbows and shoulders. He drew lanes towards the left and right at multiple levels to cover the steel surface. At one moment, he lowered and flipped the gun upwards to spray the bottom. Because the beam was short enough, he then slightly turned and raised his arms fully to spray its furthest points (Figure 4-A). He supported his left-and-right arm routine by slowly swinging his body in the same direction. On Level B, Fireproofing Worker 1 was expected to apply double layer on columns, reaching a thickness of one inch and seven-eighths for a two-hour rating, and single layer on beams, meaning a thickness of five-eighths of an inch. To ensure these numbers, he was keeping his arms’ pace and distance along with the number of passes consistent.

The order that Fireproofing Worker 1 followed in spraying was improvisational. After spraying the north-south beam, he lowered the lift. With his left arm holding the active gun towards the beam’s opposite side and his right hand on the lift’s controller, he drove further into the next section. Taking advantage of the short span, he then pivoted in a tiny rectangle formed by this first beam and the next parallel one, as well as two other east-west beams cutting across. He raised the lift back up. Now that he was facing the opposite side of the first north-south beam, he finished spraying it by swinging left and right (Figure 4-B). Flipping the gun up again, he finished spraying the bottom. Taking clockwise turns, he then finished the first east-west beam, the next north-south beam, and the next east-west beam, in this order. The next section was longer than this one. He lowered the lift, drove, and entered. This time, he started with the east-west beam running parallel to his lift (Figure 4-C). Repeating the same movements, he sprayed half of the beam within his reach, turned anti-clockwise and sprayed the north-south beam, and turned again and sprayed half of the east-west beam across. He finally drove further and finished the remaining halves of the east-west beams.

Compared to Fireproofing Worker 1 who works in a more stationary way, Fireproofing Worker 2’s supporting role demands more walks and active multi-tasking. When Fireproofing Worker 1 was working on the east-west beams further, he quickly swept the debris left under the previous ones that were sprayed (Figure 4-D). This was his first supporting activity. When it was time for Fireproofing Worker 1 to switch to the next section, I captured his second supporting activity. With verbal directions from Fireproofing Worker 1, he walked towards the west and turned off a switch. This stopped the flow of fire resistant to the gun. Switching off only allows for a short gap before “the stuff [fire resistant] will start to set up in the hose,” as I learned from Fireproofing Worker 3, and it was the same in this instance. After quickly touching base with Fireproofing Worker 3, Fireproofing Worker 1 used this gap to move to the next section. During this transition, Fireproofing Worker 2 performed his third supporting activity by collecting the gun’s hose like a fisherman would collect

their nets and clearing the way for Fireproofing Worker 1 to drive backwards (Figure 4-E). When Fireproofing Worker 1 finalized the lift's positioning, he switched on the gun.

To keep up the pace in these multiple tasks, Fireproofing Worker 2 have also developed a set of mechanics for his cleaning movements. When Fireproofing Worker 1 was working further away, he cleaned the debris under the beams he had just sprayed for a longer duration, which allowed me to observe his mechanics. One technique that I saw him use was walking straight with one relaxed arm holding the squeegee in front of him (Figure 4-F). He used this technique to clean long lanes. He walked lanes from west to east scraping the debris. In tight spots, he employed a different technique, bending down slightly and using his arms to make short quick scraping movements (Figure 4-G). Ultimately, he collected all the debris at a further spot he designated. At the end of the day, he shoveled it up into a leftover box to be pumped or transported to the ground level.

Other than continuity and pace, the task also involves mobility of the crew going between columns and beams, often in a designated portion rather than across the entire floor, unlike for cleaning and housekeeping. The workers' pace and amount of navigation in each portion are conditioned by the number and size of the components to be sprayed. On Level B, they shifted to a new portion every other day. For instance, by the time of my second observation, the workers had already advanced to the southern half. Unlike in the northern half, they worked on the lower part of the columns first. Therefore, Fireproofing Worker 1 was initially spraying on the ground. Meanwhile, Fireproofing Worker 2 covered the holes around the columns and helped maneuver the hose when Fireproofing Worker 1 switched between columns. After the first layer of coating was complete, they continued with the higher parts and the beams. For this, Fireproofing Worker 1 was back on the lift going to each section as he did during the previous observation, and Fireproofing Worker 2 was trailing him to scrape and collect the debris on the ground as well as control the switch and manage the hose.

The other significant quality is fireproofing's collaborative nature. Part of the collaboration is coordination with the other trades to arrange the open area that the task requires. This was conducted by Fireproofing Worker 3, who scheduled meetings with the Superintendent and the other foremen. Accordingly, he outlined the zones where Fireproofing Workers 1 and 2 would work and in which order they would proceed over the week. However, things on the floor are not always as smooth as planned. A delay by one trade in an area may cause shuffling, and, borrowing from Fireproofing Worker 3, that is when migrating on the floor may become analogous to "moving troops." In such cases, Fireproofing Worker 1 or 3 negotiated with other workers and improvised.

Alongside inter-trade coordination, the other part of the collaboration is the teamwork I observed on the floor between Fireproofing Worker 1 on the lift, Fireproofing Worker 2 on the ground, and Fireproofing Worker 4 at the mixer. Although each worker's role was loosely scripted, their actions were improvisational based on the flow.[4] This part of the collaboration is also where I again captured the less visible intellectual layers of a construction task. Fireproofing Worker 1's sense of viscosity and his guiding Fireproofing Worker 4 at the mixer to maintain the required thickness is one example. The day-to-day humidity affects the fire resistant mix, potentially making it too wet or too dry, as Fireproofing Worker 3 informed me, so Fireproofing Workers 1 and 4 at times had to communicate about "mak[ing] it looser, mak[ing] it tighter, add[ing] more water, back[ing] off the water" in order to "keep it [the viscosity of the mix] consistent." Besides, Fireproofing Worker 1's leading on the floor, checking proximity to adjacent tasks, and determining how to navigate an area in the most efficient way possible is another example. Fireproofing Worker 2's planning ability in

4 The part of the crew, which consisted of a "ground guy" and an "up guy," as many workers call it, was a typical collaborative unit that I would also observe in other tasks in the coming days.

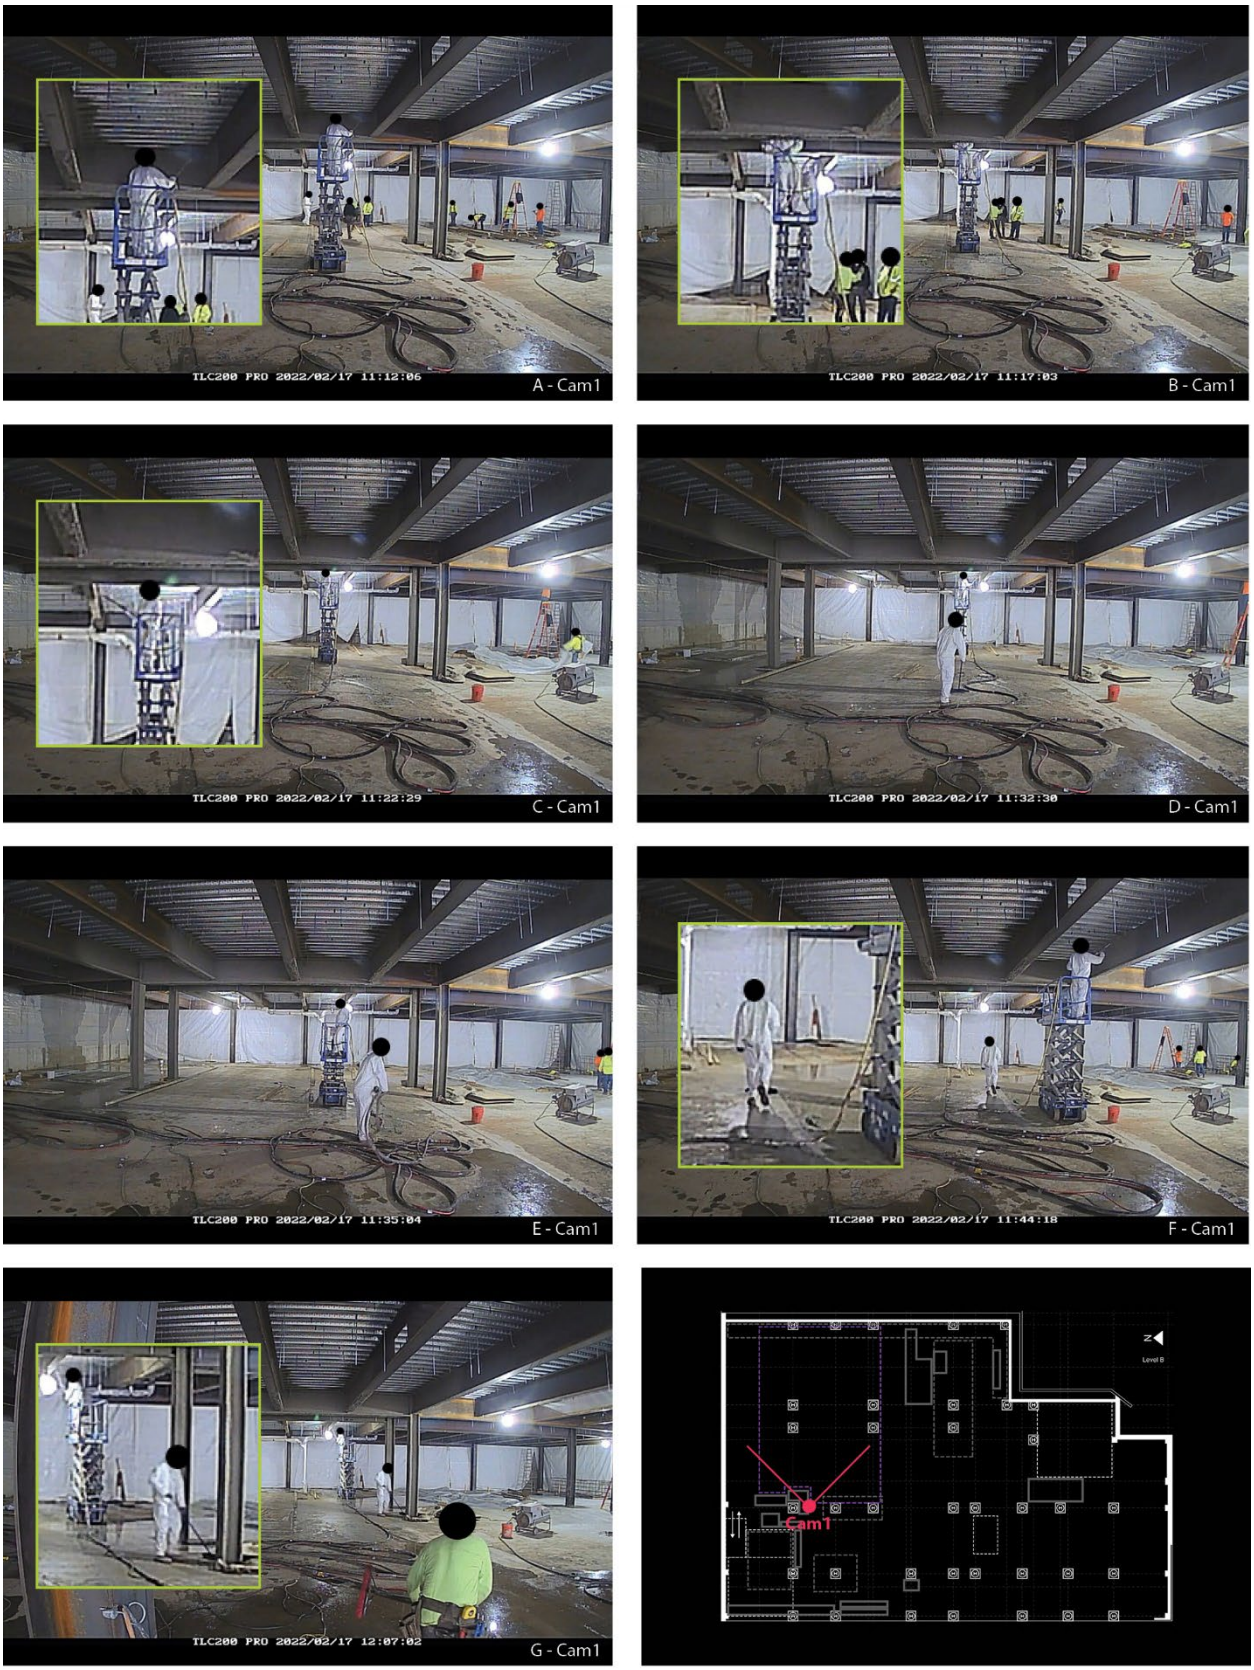


Figure 4. Key moments of Fireproofing Workers 1 and 2's workflow on Level B

cleaning and hose management is also important to highlight. Similar to Cleaning-Housekeeping Worker 1, he works with the site to find spots for debris collection and arrange the hose in a way that is both convenient for the workflow and not disruptive to others.

Rebar Tying

> *We work together all day every day, and he knows my next move, I know his next move, and we can work down a line. It's not like every move is choreographed or set in stone. (Ironworker 2)*

As the name of the task indicates, rebar tying consists in tying horizontal and vertical steel rebar units with wires to build a rebar cage, which serves as a structure to reinforce building components such as concrete floor slabs or walls. The task is performed by a specialized group of tradespeople called "ironworkers." The size of an ironworker crew is determined according to the required efficiency the workers had to deliver as well as the size of the component that they would build the cage for. The cage for a long low wall that I observed being built on Level B was executed by a pair of them, Ironworkers 1 and 2. A cage for a smaller wall could be composed of shorter bars and fewer crossings, so a single worker may handle the task. To efficiently tie a large-scale mesh for a floor concrete slab, a larger crew may be required.

Rebar tying took place on the first day of fireproofing on Level B's eastern half. South of where Fireproofing Worker 1 was traversing, a pair of MEP workers were checking their hardware on the ceiling, and this same area also featured Ironworkers 1 and 2 in action (Figure 5). This pair were making rounds between the long rebars that Cleaning-Housekeeping Worker 1 stacked the day before and a line of short rebars sticking out of the floor vertically. By the time I finished observing the fireproofing workers, part of the rebar cage the pair was expected to build was visible. The height of the vertical rebars suggested that the cage would be for a low wall, and the length of the artifact, partially invisible behind a tarp, indicated that it would circumscribe the entire eastern edge. I began watching the workers checking the sizes of pieces, bent down on each end of the rebar stack. They lifted a bundle they selected and carried it all the way to the eastern edge. They carefully lowered the bundle and dropped the pieces. When I walked closer, I saw the rebars sandwiched between two parallel rows of verticals. The workers made another trip to the stack, walked back to where they had dropped the bundle with another rebar in hand. Their tying flow was about to begin.

While less physical compared to some other building tasks, rebar tying is highly dexterous as one could observe by watching the ironworkers' hands. Ironworkers multi-task with fingers, holding the rebar and a piece of wire in one hand, and knot the wire around a horizontal-vertical crossing with a small tool called the "plier" with the other. This artful move is repetitive and, when repeated every few seconds, looks mechanical. Yet, every repetition is masterfully executed with control in handling the rebar and agility in tying it as Ironworkers 1 and 2 demonstrated on the site. The two workers positioned the last rebar they brought horizontally at the highest level along the row at the back. While Ironworker 1 was holding it at approximately the one-fourth mark on one end, Ironworker 2 was holding it at its opposite end (Figure 5-A). I realized that this was meant to prevent the twenty-foot-long bar from sagging in the middle. In approximately one second, both workers started to tie the horizontal bar to the vertical one closest to them. In an interview, Ironworker 2 described the details of this quick, skillful tying action. "You got to hold the bar in place," he explained, showing his thumb and index fingers holding an imaginary rebar. "As you're holding it, you got a piece of tie wire in this hand already," he continued, holding a piece of wire with his middle, ring, and pinkie. With a plier in his right hand, he rolled the wire around the imaginary rebar as the final step.

Because of the long size of rebars, Ironworkers 1 and 2 had to work in tight coordination in carrying, positioning, and tying rebars. It was visible that while these workers tied the rebar, the rebar tied them together. One worker's movements occurred according to the other's. The tying order of intersections was not scripted but rather improvisational: whoever was faster would see a part sagging and tie it. Ironworker 1 tied his first intersection faster, which made the rebar stable on his end and freed his hands. He moved to a new intersection in the middle and tied it by using his left hand to help his right hand operate the plier (Figure 5-B). He then moved closer to Ironworker 2 to tie another one in a repetitive manner, walked back to the opposite end, and tied three more (Figure 5-C). Meanwhile, Ironworker 2 tied a few on the portion where he began. After these, he moved up to tie one more, and while he did so, Ironworker 1 moved closer to him and tied his last one. Ironworker 2 took one step to his left and tied his last one as well while Ironworker 1 checked the intersections he executed earlier (Figure 5-D). This swift, coordinated movement that I watched for approximately three minutes was similar to a dance. Every next move was partially dictated by the blueprint of the cage specifying the number of ties along a particular rebar length and partially improvised by the workers trying to be time efficient.[5]

The rebar also tied the two workers through communication as I noticed. Part of this communication was verbal, through what Ironworker 2 called "our own phrases." An example is "hold that," which he used to stop Ironworker 1 while carrying or tying. When positioning or coordinating movements, both workers employed short instructions such as "to me," "to you," or "my way." The workers used verbal communication for motivation as well. As Ironworker 2 explained, this was intended to "work at a pace" or "lift up spirits." However, a big part of the workers' communication during tying was silent. "The more you work with someone, the more you learn them," Ironworker 2 said, referring primarily to their movements. "If one of you gets switched, would it [communication] work that way?" I asked. "Absolutely no," he replied. His answer indicated that while knowledge of the task is unified through apprenticeship, communication between workers during the task is customized.

Similar to fireproofing, the conspicuous quality of the ironworkers' tying flow I captured was its continuity. Its distinctive aspect was that both workers were highly mobile and operated at a visibly faster pace within this continuous flow. After the first rebar was in place, Ironworker 1 quickly bent down and picked up the tip of the second one, which was shorter and required fewer ties. Ironworker 2 walked to the other end, bent down, and lifted the tip to the required height, which he found based on a mark on the vertical bars (Figure 5-E).[6] Once the rebar was positioned at the mark, Ironworker 1 quickly tied his first intersection. Ironworker 2 moved down about one-third of the bar and tied his first a second later. Ironworker 1 was ahead again. After stabilizing the end, he shifted closer to Ironworker 2 and tied his second intersection. While he moved on to lift the third rebar, Ironworker 2 moved back to the end where he started and tied the last couple of intersections. Then, he joined Ironworker 1, and they positioned the third rebar to the proper height (Figure 5-F). Through similar coordinated movements, they finished tying the intersections. Subsequently, the fourth bar and the fifth bar, which were close to the ground, were lifted and completed in the same way. Whereas standing or, at times, bending down slightly was enough during the first and second rebars, with the last three bars, Ironworker 1 bent down and squatted, and Ironworker 2 knelt down at the intersections (Figure 5-G). On average, the workers moved linearly along the cage every few

5 Aside from the location and rebar size, the blueprint specified the spacing between rebars and, accordingly, whether "the bars need to be tied hundred percent," meaning "that every intersection those bars make get a tie," as Ironworker 2 explained. "If we have large spacing [between vertical rebars], we'll tie every bar," he stated and continued: "if they are very close together, we'll gauge what we can with our experience." Hence, the workers did not always need to tie every intersection, as I also observed on the site. However, with experience, they knew which ones were critical for the cage to hold intact.

6 Transferring these points from the blueprint to rebars is what constitutes the layout work in rebar tying. It is often done by tape-measuring to mark the points with paint on rebars.



Figure 5. Key moments of Ironworkers 1 and 2's workflow on Level B

seconds and jointly tied more than ten intersections in approximately one minute. They continued this as a cycle, which made the task seem mechanical.

The workers' continuity and pace came partially from the pressure of tying quantities of intersections. After finishing the segment I described, the two workers then did four more rounds between two partially completed segments and the rebar stack. The first round was to add a long rebar to Segment 1 running north-south. This rebar was added to a point where it formed a corner with Segment 2 running east-west. The second round was to make shorter rebars for Segment 2, which would connect to the longer rebars extending from Segment 1. For this, Ironworker 1 picked up a power tool used for cutting rebars, called the "angle grinder." Wearing a welding shield, he trimmed a piece, which he then passed to Ironworker 2 to tie. While Ironworker 2 worked on the short pieces of Segment 2, Ironworker 1 did the third round, carried a bundle of long rebars, and began to tie them to Segment 1. He had to do some more cutting in the middle of this flow. Finally, in the fourth round, the pair finished the back side of Segment 1 and switched to a new zone. In another forty minutes, they fully completed the back sides of a long and a short segment running north-south; at that point, they had tied all the rebars on all the back sides. Finishing the back sides first was part of their plan, as Ironworker 2 explained during the interview, because tying the rebars on the front first would block their way to the back. The pair finalized the cage by completing all remaining front sides. They stored the leftover rebars on the stack, covered it with a tarp, and called it a day.

To keep up this pace, Ironworkers 1 and 2 carried only a small number of items, including pliers, wires, and a knife, in their tool belt harnesses. Stationing other tools they use only when needed, such as the angle grinder for rebar cutting or the tape measure for the layout, was part of weight management. The only heavy things they carried were rebars, which made their task less physically demanding than many others. Depending on their preference and congestion on the floor, they would hold and carry them with their arms hanging or on their shoulders.

Also similar to fireproofing, rebar tying is a highly collaborative task, but with idiosyncrasies. The collaboration on Level B began before the period when I observed the workers, with engineers passing the blueprints to Ironworker 1, who was the foreman leading this job. The collaboration between the workers started with the drawing by getting on the same page about the rebar dimensions and spans. The requirements of the long low wall cage they built were relatively straightforward, but those for a load-bearing wall or concrete slab reinforcement could be more intricate, as I discussed with Ironworker 2. The collaboration continued with planning the order of segments to complete. The plan was contingent on the location and movement of other crews and, therefore, determined in coordination with the Superintendent and other trades. For instance, the workers started the segments further from the fireproofing crew and came back to those under the fireproofed beams later. This part, in which they reviewed the requirements and planned their flow, constitute their intellectual labor. Finally, the collaboration continued as tight coordination during the tying flow that I described. In this case, the workers performed mostly identical roles of carrying, positioning, and tying rebars, which differs from fireproofing, where roles were complementary.

Carpentry Formwork

> *I feed him [co-worker] panels, clamps. I bring tools. That's my job, why I'm here. (Carpenter 2)*

Carpentry formwork is the task of building a mold called "formwork," which is used to hold cast-in-place concrete until it dries and forms building components such as slabs, walls, or columns. While a

formwork could be made of different materials for different components, on Level B, it was created by installing aluminum panels along both sides of the rebar cage the ironworkers built (for the low wall on the eastern edge). Depending on the size of formwork to be built, the size of the crew may vary. For the low wall on Level B, the crew included two carpenters, Carpenters 1 and 2.

The day I observed carpentry formwork, Level B was starting to be more congested with MEP materials, including gang boxes, tool carts, hangers, pipes, and conduits stacked in different zones. An air handler had been placed on a raised concrete base called the “housekeeping pad” the day before. This large white box-shaped equipment now divided a future mechanical room on the northeast from the rest of the floor. In addition to more materials, there were more trades. I saw the fireproofing crew continuing with the southern portion of the eastern half. Steam fitters and sheet metal workers were starting to fix the air handler in its location. On the west and east sides of the handler were plumbers and electricians, respectively, preparing for a day of ceiling installations. Two workers whom I thought were doing the formwork panel installation were situated at the northeast corner (Figure 6). “This is carpentry formwork,” said one of them when I asked him to confirm.

In the workflow that I observed for the next hours, Carpenter 1 was the lead worker carrying out the panel installation, and Carpenter 2 was performing a supporting role, setting temporary panel stations, transferring panels, and delivering tools and hardware. The worker who answered my question was Carpenter 1. He was standing in a roughly five-foot-wide space between the back of the cage and a retaining wall climbing up to the ground level. The ground underneath him was lower than Level B’s deck, and, therefore, the top of the cage that would normally be around his chest was aligned with his helmet. Three panels were already fixed on this side of the cage. The supporting worker, Carpenter 2, was on the deck on the other side.

Carpenter 2 in the supporting role was noticeably more mobile, taking more frequent, longer, and faster walks between multiple destinations, including a workbench where the workers kept their tools and hardware, the main station where they stored panels, and temporary panel stations he set in proximity to points of installation. Carpenter 2’s work was also more physical as he lifted, carried, and maneuvered panels repetitively. Following a verbal cue, I saw him walking up to a hardware pile located near the workbench at the starting point of the cage. He picked up a few pieces, walked down to Carpenter 1’s location, and threw them to his side (Figure 6-A). A few steps behind him was a temporary panel station, which contained the panels he had stacked on a pallet. As I noticed earlier, the main panel station lay further south, perpendicular to the cage. Carpenter 2 turned to the temporary stack, lifted one of the panels to his chest, and rotated it with his hands. By resting the panel on the cage, he flipped it over to Carpenter 1’s side (Figure 6-B). He extended his arms through the cage and pushed the panel to lean on the retaining wall. This was the third panel he queued up for Carpenter 1 to pick and place. He then walked back to the hardware pile and transferred a nearby red cylindrical tank with a hose back to where he was. With the same sequence of moves, he moved a fourth panel to the other side. After briefly touching base with Carpenter 1, he started spraying all four in the queue with a liquid coming out of the hose. “If I don’t [spray], concrete will stick [to panels],” he explained. These were his panel preparation activities.

In contrast, Carpenter 1 was performing the dexterous part of the labor, aligning panels, joining them with a special hardware called the “clamp,” and stabilizing them with a hardware called the “bracing.” Thus, the workers’ collaboration involved complementary roles, which separated the dexterous and physical parts of formwork. While Carpenter 2 was constantly moving to perform the panel preparation activities, Carpenter 1 remained fixed on his ground illuminated by the beams of a tower light. He picked up one of the panels Carpenter 2 had queued up for him and lined it up next to

another at the edge of the cage (Figure 6-C). The new panel was thus adjoined to the row. To fix the free-standing panel, he picked up one of the clamps Carpenter 2 had passed to his side. He placed the clamp at a high point where the two panels aligned on their sides. With a power tool called the "impact gun," he fastened the bolts on the clamp. The fastened clamp held the new panel and its neighbor together. He would do this multiple times throughout the workflow.

Installing a unit of material over and over is inherently repetitive work with steps that are more or less scripted. This would be a good summary of both the ironworkers' tying and the carpenters' formwork tasks. However, compared to the ironworkers' swift, serial tying move, Carpenter 1's fastening cycle required carefully positioning a heavier material as well as handling more hardware and tools. Each of his predefined steps took longer than ironworkers' one-step tying move executed in a few seconds. Consequently, each panel took longer to fix than a rebar to be attached. This expanded time span for each unit made Carpenter 1's repetitive flow look less mechanical, with more spontaneous moves in the mix. It also made Carpenter 1's overall pace of movement along the cage slower. While ironworkers moved the span of a rebar every few seconds, he moved the span of a panel in the time it took to complete the installation, which was often two to three minutes per panel.

Different from installation, Carpenter 2's supporting activities did not necessarily consist of any predefined steps but rather of what he had found convenient over time. Still, the order between them could be tied to Carpenter 1's flow as further parts of the workflow showed. Continuing his back-and-forth traversal, Carpenter 2 shifted the tank to the next segment to which Carpenter 1 would soon arrive. He then walked back to the temporary station and pulled the pallet jack carrying the last two remaining panels to the same segment (Figure 6-D). With the movement routine I described, he flipped and passed the first one over to Carpenter 1's side. For a few seconds, he walked back and removed the wood beams that worked as the pallet. He then returned to the segment, moved the pallet jack to open up space, and passed the second panel. He sprayed both panels. Meanwhile, Carpenter 1 placed his second panel and started fixing the clamp that would connect it to the first (Figure 6-E). Each time Carpenter 2 noticed Carpenter 1's need or Carpenter 1 asked for hardware, delivery became the priority over his panel-related activities. For instance, after spraying, Carpenter 2 came back to his area, bent down into a container of clamps, picked up a few, and threw them to his side. He did a few more rounds and laid more clamps along the path Carpenter 1 would follow.

Nonetheless, Carpenter 2's panel-related activities were mostly independent but remained somewhat tied to Carpenter 1's flow given that he had to remain ahead to line up the panels. Once Carpenter 1 had fixed the second panel to the first, he moved on to adjoin the third to the second. At that time, Carpenter 2 was preparing the next temporary panel station. As the first step, he pulled the pallet jack to the main station in the south of where they worked, containing eight stacks of panels on pallets. Then, he pushed the jack under one of these and loaded it with a new stack. Finally, he pulled the loaded jack and parked it where Carpenter 1 would reach a few panels later (Figure 6-F). As the new panels arrived, Carpenter 1 was done fixing the third panel and began working on the fourth. At this moment, Carpenter 2 walked a short distance to pick up a cutter from the workbench. Together with the clamp container, this long wood bench made this area into the main tool and hardware station, holding the less-used lasers, level, and angle grinder as well as bolts in a container. With the cutter, Carpenter 2 untied the straps holding the stack.

As Carpenter 2 began to transfer the panels to the other side of the cage, Carpenter 1 fastened the bolts on the clamp between the fourth and fifth panels. With the fifth panel in, a segment spanning the length of a column was covered (Figure 6-G). This was when Carpenter 1 began stabilizing the formwork. He attached the end of a long steel rod called the "bracing" to one of the clamps. He



Figure 6. Key moments of Carpenters 1 and 2's workflow on Level B

attached a second clamp to the rod's other end and adjusted its size to extend to the retention wall. In the same way, he installed a bracing as a support to each panel couple. As he was transitioning into the next segment, all panels, clamps, and bracings he needed were lined up. He collected the power tools he kept nearby and carried them with him to the next segment. I continued observing this routine of Carpenter 1 installing the panels and Carpenter 2 moving between stations and bringing him hardware and tools as they almost finished the entire back side that day. With a third worker's help, they fully finalized the remainder and the entire front side the next day.

Similar to fireproofing and rebar tying, carpentry formwork is a collaborative task. Part of the task, including the arrangement of the main panel station and workbench as well as the side of the cage and the direction in which to proceed, was planned in advance in coordination with the Superintendent and other trades. Ideally, panels, tools, and hardware should conveniently be accessible from different points along the cage, and the direction should fit with other trades' plans. The task also involved on-the-fly planning, such as setting up the temporary panel station based on needs and available space. Lastly, the sharing of roles between the workers I noted was also a part of the plan. The roles were complementary, similar to those involved in fireproofing.

3.3 Key Insights

Fifteen weeks at the New Scaife Hall site produced key insights, which I synthesize in three sets and present in this section. The first set explains the way work is conducted on a construction floor, including the spatial arrangement, movement, and interactions. The second set is drawn from the task descriptions, and details the qualities of tasks to advance the generic descriptions of construction work. The last one presents workers' perspectives on robots, which they expressed during the interviews.

These insights and their detailed descriptions in this section convey the ReAC group's common understanding of the site conditions and interactions, attention to tasks' specific qualities and the specific ways they are enacted—rather than their general character or commonalities—, and incorporation of workers' voice into our conception of systems, all of which I expand on in the next chapter. These shaped the way we analyzed tasks to be deeper and detailed, which then led to our conception and demonstration of an alternative supporting role for a robot, as I previously noted in this chapter and expand further in the next. The understanding of the site conditions and interactions described in this section also informed the technical development of the final prototype, which I specify in the next chapter.

3.3.1 Way of Working: Zones and Interactions

The first key insight is that workers organize materials spatially in designated zones, and these, more often than human traffic, impede movement on a construction floor. The New Scaife Hall site did not have enough outdoor space for storage, as Plumber 4 pointed out. Thus, trades had to store every batch of their material on the floors. Each time a trade received a batch on Level B, workers stationed it on designated "lay down areas," as they called them. These zones began to appear on Level B with steam fitters', plumbers', electricians', and carpenters' initial material deliveries, ushering in a new order of congestion (Figure 7). The floor, which was wide open after cleaning-housekeeping and kept that way for another week during fireproofing, became more crowded with these materials. With the addition of fire sprinkler fitters', sheet metal workers', framers', and drywall workers' stuff, the material zones propagated (Figure 7). With new batches added each week, the zones grew in size and number, forming a denser tapestry in a significant area. Each zone,

depending on its owner trade, contained large materials including pipes, ducts, tracks, studs, or drywall panels on wheeled carts, smaller materials such as rods, hanger heads, and assembly hardware in containers, and various tools in gang boxes, rolling scaffolds, and carts. Alongside the interior wall frames, these items were the largest obstacles that constrained the number and width of the paths for worker, lift, and material circulation.

Contrary to the common description of sites as unstructured, these zones show that there is a degree of order in how a floor is configured. Workers arrange them in a premeditated and coordinated manner based on a couple of principles. One principle was to put materials as far as possible from busy areas. For instance, steamfitters' stacks were mainly placed on the north and west sides of Level B (Figure 7). Steam Fitter 1 explained that these areas would receive the smallest amount of ceiling work, so the stacks would not "interfer[e] with other people." For the same reason, the fire sprinkler fitters kept their stacks on the southwest corner (Figure 7). "Laying them out in a way, so that you move them only once to install later on," Fire Sprinkler Fitter 1 said to summarize their strategy. Similarly, framers distributed their tracks and studs to the middle, the south, and the nearby drone test arena (Figure 7). As I previously described in detail, drywall workers stored their panels in the southern portion and around the arena (Figure 7). The other principle was to put the stacks as close as possible to the area of installation, keeping the materials' weight in mind. Conduits are the lightest of all MEP materials, so electricians stored most of them in the northwest without concern for carrying distance (Figure 7). Conversely, Plumber 4 wanted to keep their main stack a short distance to the middle and southwest sections, where their lines were heavily concentrated (Figure 7). Because sheet metal ducts can block space, keeping them close is not always an option: "In an ideal [world], if I have the floor to myself," said Sheet Metal Worker 1 in a cynical tone. Nevertheless, the crew still managed to put most of their pieces within the eastern half, not too far from their installation areas (Figure 7).

Similar to materials, work is also spatially organized on a construction floor. As I highlighted in the task descriptions, the Superintendent and trades plan the positioning of tasks because each creates constraints for the rest, either by leaving material debris, generating lift traffic, or requiring a route along certain components such as beams, rebar cages, hangers, or frames. According to this positioning, trades perform tasks in virtually bound zones. Each zone was typically run by a unit consisting of two ground workers or an up worker coupled with a ground worker carrying out a task. Alternatively, a unit working in a zone may include two up workers and a ground worker. Starting from my second week, the number of these zones increased with the propagation of material zones, which made the floor congested (Figure 7). On average, twenty to thirty-five workers along with their tool carts, temporary material stations, and lifts in work zones, together with their material stacks, took up almost the entire floor space. Hence, squinting one's eyes from above, construction is spatially a dance between work and material zones on a floor, not strictly choreographed but loosely outlined by coordination between trades. The moves are still not perfect; Fire Sprinkler Fitter 1 expressed his frustration in this regard through an example: "you might show up [for a task]… and then you find out that there [are] ten piles of drywall in that area." Sometimes, this occurs due to the "screw it" attitude of workers, as he calls it, which frees them from accountability and causes such conflicts. Hence, within the current organization of a floor, there is no perfect way to cater equally to oneself and to others.

Another key insight on the way of working is that workers' movements did not necessarily concern the entire floor but were rather confined within work zones. As I stated in the task descriptions, the tasks were mobile to different extents. Some, such as cleaning-housekeeping and fireproofing, required traversing a wider area and changing locations more rapidly in a day. Others that

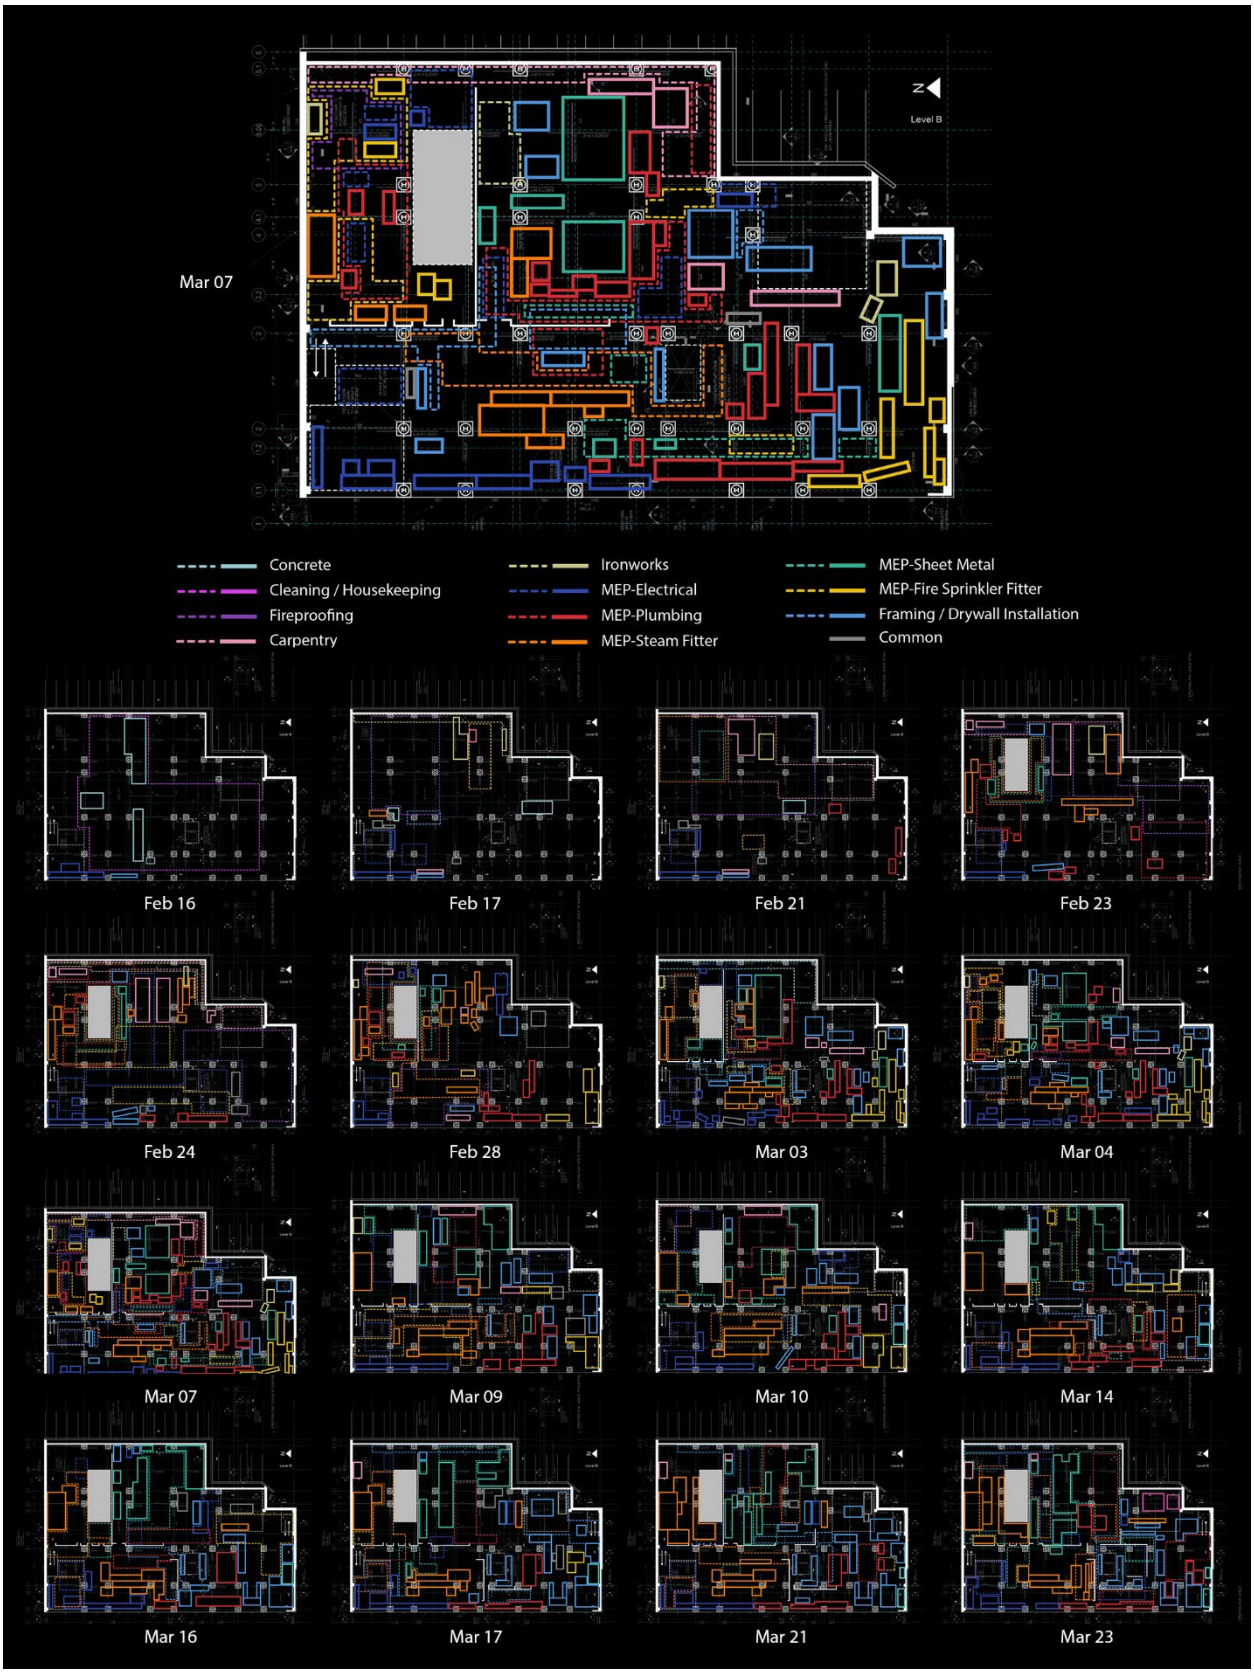


Figure 7. Work zones (dashed boundary) and material zones (solid boundary) on Level B

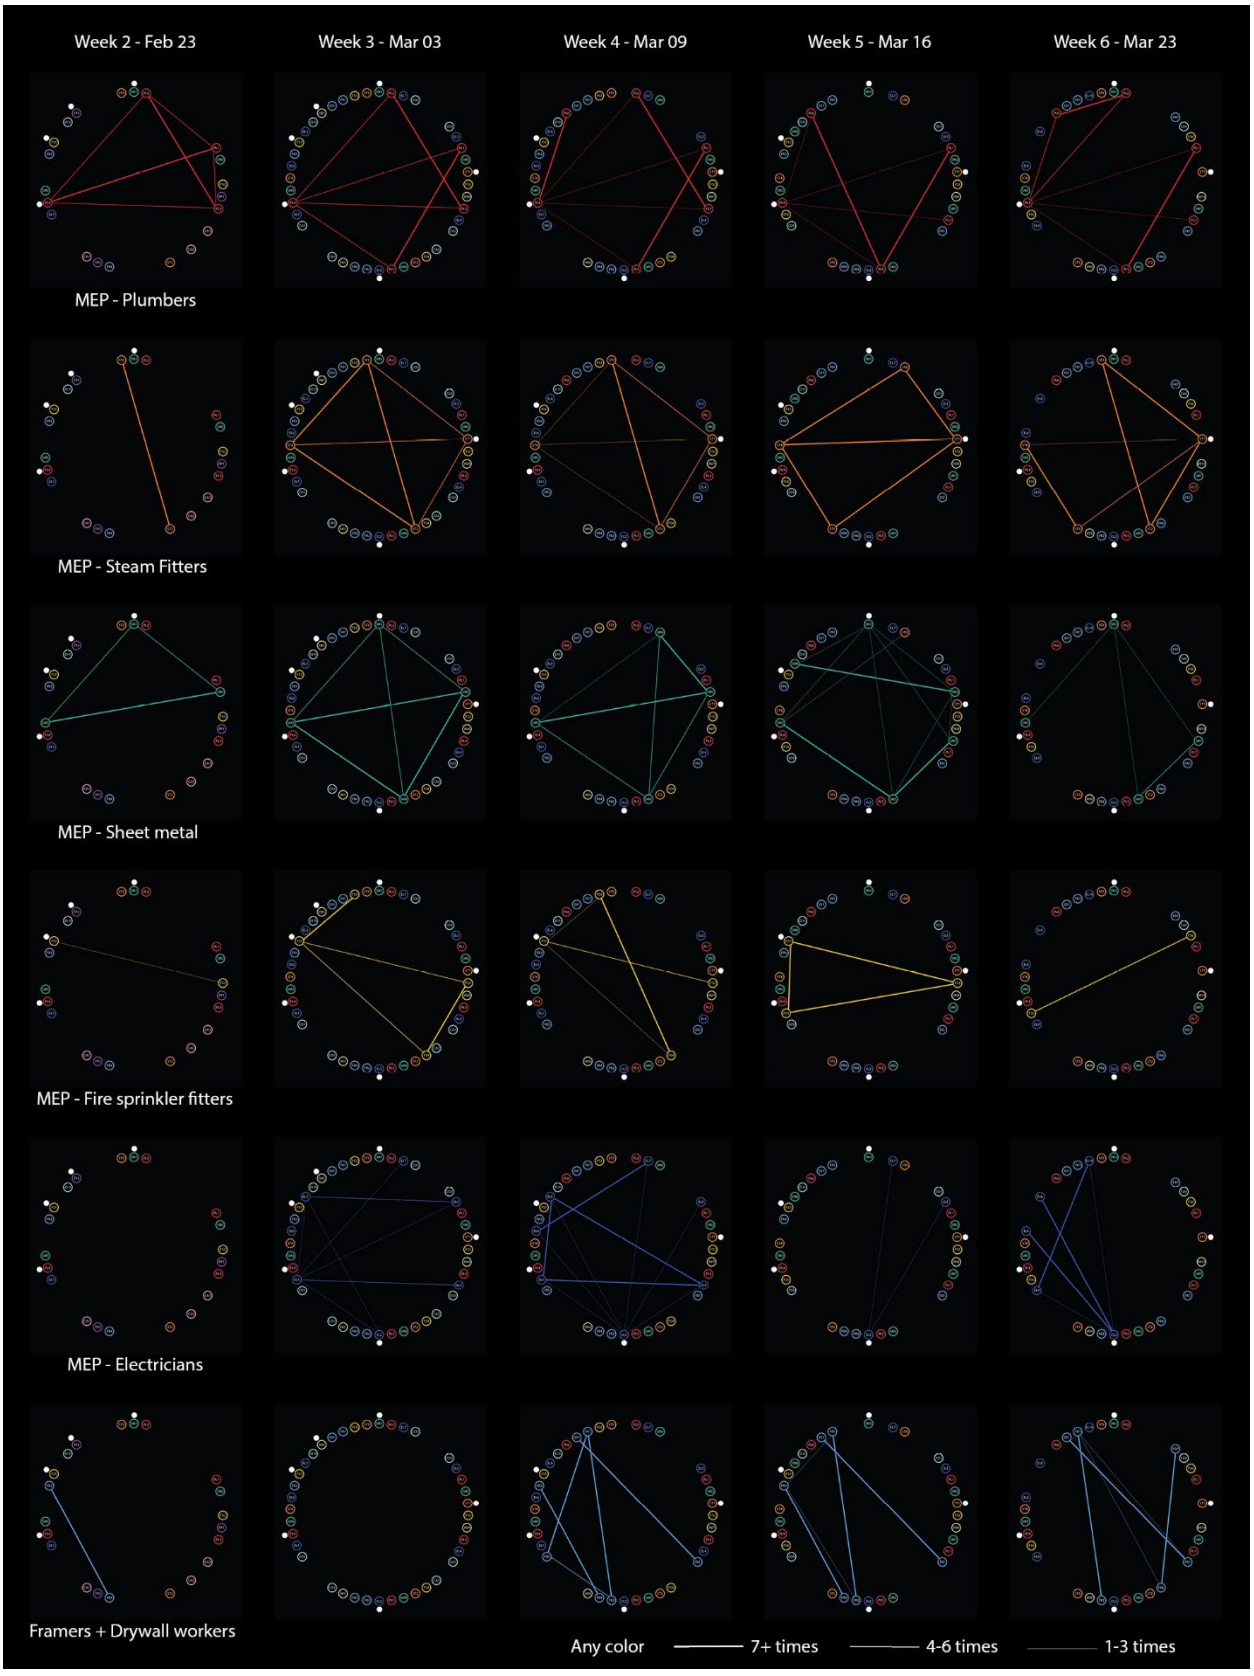


Figure 8. Daily work interactions in trades, showing more interactions among workers in a zone than across zones

proceeded as fast as the assembly, such as rebar tying, carpentry formwork, MEP ceiling installations, framing, and drywall installation, entailed movement along a rebar cage, frame, or ceiling portion and between that artifact and a material station. In all cases, workers tried not to exceed their zone as they carried out their tasks. For instance, I observed the fireproofing crew operating in an approximately seventy-five-by-forty-six-foot (twenty-three-by-fourteen-meter) zone. The ironworkers were tying within a linear one-hundred-and-nine-by-nine-foot (thirty-three-by-three-meter) zone, when I was shadowing them. This is firstly because units had to maintain distance to maximize safety and secondly because they had to keep their work separate to minimize conflicts. An unwritten rule was to respect the zones of others, so movement between work zones was minimal, occurring when a unit completed its activities in one zone and migrated to another. Depending on the task, units worked in two to three zones a day. This also happened unavoidably when work zones overlapped, which was common since the organization was never perfect, as I highlighted.

The last key insight in this set is that collaboration in construction was not as highly interactive between every single worker as one might think. Interactions were more common within a unit or a trade than between trades. Interactions between units in a zone, such as the ones I portrayed in fireproofing, carpentry formwork, MEP ceiling installations, framing, and drywall installation, were the most frequent and collaborative. These interactions were for planning the day, communicating drawings, assembling large materials, or passing tools and materials. On the site, they usually

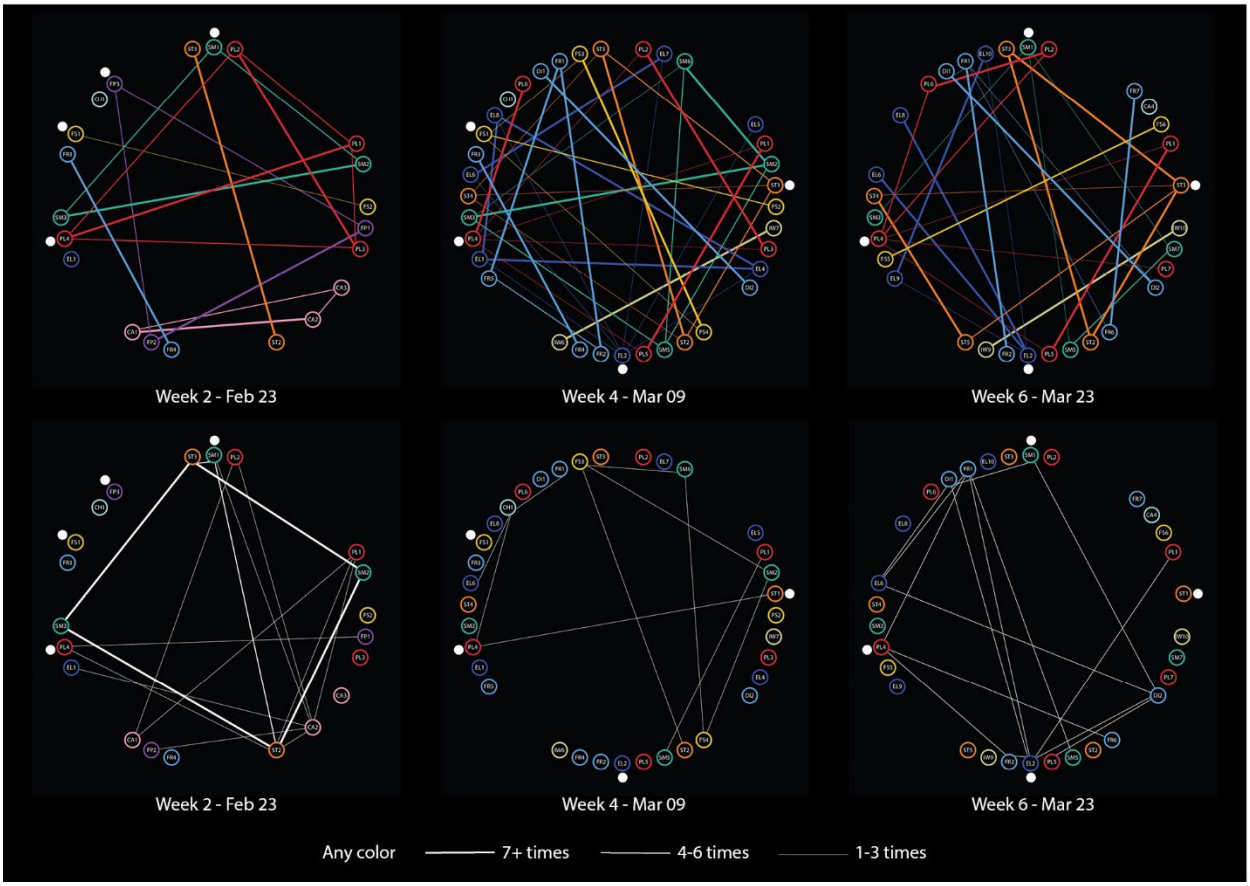


Figure 9. Daily work interactions within trades (top row) and between trades (bottom row), showing more interactions within trades than across trades

occurred more than seven times a day between each pair within a unit (Figure 8). Units of a trade in different zones were independent, but workers across these units may still touch base one to six times a day for planning, consulting each other, or resolving lift traffic or conflicts (Figure 8).

Conversely, interactions across trades were limited and coordinative in nature (Figure 9). They were mainly triggered by three reasons: resolving conflicts, negotiating work zones, or checking for a floor-wide shared instrument. Individual workers or units from different trades interacted when there was a conflict, for instance, when installing a steam pipe first would hinder the fitting of a duct. Workers then revisited the drawing, as Sheet Metal Worker 1 explained, and with "quick imagination," reasoning such as "I can't lift my guy there and dance around the steam fitter pipe" led to a coordination talk. Workers from different trades also interacted when they had to negotiate a work zone. "Hey… I want to do this in two days, so can you get over there before that?" was Sheet Metal Worker 1's pitch to the plumbers in a hypothetical scenario where his crew gave two days to the plumbers in a work zone that they wanted to work on afterward. Workers across trades also interacted to check for a shared tool on the floor, for instance, the framers asked when the carpenters would be finished with the pallet jack. Overall, there was not much give-and-take between trades in terms of collaboration, tools, or materials.

3.3.2 Qualities of Work: Grounding the Generic Descriptions

Terms such as "unstructured," "physical," "(un)skilled," "repetitive," "mobile," or "collaborative" indicate a common blanket character for construction work. However, they do not necessarily convey how each task embodies these qualities. What specific actions in a task create the impression of these qualities? For instance, if a task is mobile, in what specific way is mobility really enacted? A key insight that I gained by observing construction tasks is that they do not always bear these qualities as absolutely as portrayed. To put it differently, there are nuances between how construction is depicted and what one sees on a site.

First, certain qualities of tasks are not immediately visible. They become so when one is on a site and listens to workers. For example, the arrangement of tasks may look messy and unstructured, but it is not unorderly, as I showed with the planned organization of work zones. In addition, workers who have to cross wider areas (as in cleaning-housekeeping and fireproofing), advance as they assemble (in rebar tying or formwork), or work on lifts in crowded zones (in MEP ceiling installations, framing, and drywall installation) internally plan the order of their daily activities. Execution certainly involves many improvisations. Further, workers who need to set up material stations for pipes, tracks, or panels (as in formwork, MEP ceiling installations, framing, and drywall installation) organize their work zones. Similarly, construction tasks are undoubtedly physical, but each one also involves intellectual layers. As I showed, even in a seemingly simple, heavily physical cleaning-housekeeping task, one has to take in the coordination information to devise a plan. Other than planning, fireproofing workers' knowledge of material viscosity, MEP workers' knowledge of hardware, tools, techniques, and technical specifications, and the layout activities in MEP, framing, and drywall installation are some other examples that I observed.

Second, each task or activity that constitutes a task embodies each general quality to a different extent. Not every task is heavily physical or necessitates trained skills, craftsmanship, or dexterity. For instance, cleaning-housekeeping activities require physicality rather than craftsmanship and dexterity. Conversely, rebar tying is much less physical but impossible to perform without dexterity. This is also the case of the different activities that constitute a single task: some are more skill invested, while others rely less on expert skills and more on intuitive physical movements. This

difference can be observed between spraying and cleaning in fireproofing, or between supporting activities and installation in formwork. Similarly, not every task or activity is repetitive, and some are perceived as more mechanical than others. Tasks such as cleaning-housekeeping and activities such as cleaning in fireproofing and supporting in formwork are unscripted and not based on repeating a set of unit moves, as I highlighted. Conversely, spraying in fireproofing, rebar tying, or nailing tracks in framing relies on routines of mechanical moves. Still, these activities are not free of improvisations. When the pace of assembly is slower, like in formwork, MEP ceiling installation, and drywall installation, repetitive installation is perceived as less mechanical.

Further, tasks are mobile but involve different degrees and paces dictated by the materials. As I highlighted in fireproofing, the fire resistant dictates the continuity and pace of workers. In rebar tying, it is the quantity that influences workers similarly. In formwork, MEP ceiling installations, framing, and drywall installation, the sizes of materials set the pace of assembly and, consequently, how often workers move within or across zones. Other than the materials, roles also determine the degree of mobility. In rebar tying, where workers' roles are equal, both are equally mobile; in formwork, where roles are complementary, the supporting worker is more mobile than the lead worker. Similarly, ground and up roles in fireproofing, MEP ceiling installations, framing, and drywall installation involve different levels of mobility.

Lastly, tasks involve different degrees of collaboration and communication between workers, also conditioned by materials and roles. Installing large materials, such as rebars, pipes, or studs often leads to tight collaboration, as observed in rebar tying, MEP ceiling installations, and framing. Material preparation and feeding is an important part of collaboration in formwork, MEP ceiling installations, framing, and drywall installation. Degrees and modalities of communication vary among tasks as well. Examples are the short verbal prompts between workers in rebar tying, workers asking for tools and hardware in formwork, and the communication of cutting dimensions between workers in drywall installation. In addition, silent communication occurred mostly by checking on co-workers' moves, such as a supporting worker checking tools and hardware in formwork and workers handling tall materials such as rebar in tying, pipes in MEP ceiling installations, and studs in framing in coordination.

3.3.3 Worker Perspectives on Robots: Mixed Feelings

One aim of the New Scaife Hall study was to hear workers' perspectives on robotic technologies. First thoughts surfaced in their jokes, such as "a robot should do my job" or "work harder or robots are coming," which they commonly shot at me once we built rapport. More serious, deeper thoughts were expressed during the interviews. I asked for their general reflections on robots and the ongoing quest to integrate these technologies into construction in the United States. I asked them to project into the future and share their thoughts about the role of these technologies. Specifically, I tried to get them to tell whether and where in construction they saw these technologies as useful or supportive or, conversely, unnecessary or redundant.

The situated nature of tasks: The workers I interviewed had a well-grounded view of their tasks as a subject of robotics. A salient point they commonly made concerns the potential gap between the affordances of systems and the complex, impromptu nature of tasks. Reflecting on that gap, a few workers (Sheet Metal Worker 1, Plumber 4, and Fire Sprinkler Fitter 1) began to align with anthropologist Lucy Suchman's (1987, 2007) account highlighting the discrepancy between the "situated nature" of human action and the design of intelligent systems based on its misconception as "planned action." There are plans and schedules in construction, as I previously mentioned, but,

to put it in Suchman's (1987) words, they do not "determine the actual course of situated action" on the floor. With the same awareness, these workers expressed their doubts on the viability of systems built on a similar misconception of action in construction:

> *You got a tough job [in developing a robotic system for construction] because, every day we do the same s***, but every time it's different. You have to do it differently. (Sheet Metal Worker 1)*

> *for all those things [technologies] to work and work well, it has to be planned… and the thing about construction is things are constantly changing… you could put your faith into something else, but then it's gonna change… (Plumber 4)*

> *I probably could [have robot support], if it was a standard everyday procedure, but, through the fact that nothing is the same day to day, that will be hard to do. (Fire Sprinkler Fitter 1)*

These reflections show that workers preconceived robots as technologies that would by default demand a predictable, scripted order in construction, rather than technologies that could be designed and integrated into tasks in alternative ways. Based on this view, they considered robotic technologies challenging to adapt to their complex workflows, which are at times scripted but almost always filled with improvisation. These reflections also go beyond the dominant characterization of construction tasks as repetitive and instead underscore their contingent nature.

Skills and deskilling[7]*:* Based on their general knowledge of robots combined with their expert knowledge of tasks, the workers made informative points regarding which skills may or may not easily be performed with these technologies. For instance, Drywall Worker 1 questioned whether "judgment," in particular thinking ahead in a task to accommodate a following task's needs, was a skill that could be accomplished with a system. In a thoughtful reflection, he compared a repetitive-mechanical aspect of drywall installation, screwing, to an intellectual aspect, the ability to see drywall installation and drywall finishing in a coordinated sequence. He foregrounded the latter as a challenge for any automation effort:

> *Just 'cause so much of it [drywall installation] is judgment… I put screws on certain spots so that the [drywall] finishers don't have to bend over all the time… And I know, just based on their taste, if I should use, if I can put a tapered edge into this door jam, or if I can't, or if the [finishers] are gonna get mad at me for that… The movements themselves, like screwing, putting in screws, is definitely [doable with systems]… but… there's a lot of judgment calls." (Drywall Worker 1)*

Plumber 4 had a slightly different take. According to him, even the repetitive-mechanical aspects of tasks are part of training and expertise, which grants tradespeople their current value in the market. Thus, he was concerned about potential deskilling accompanying robotic technologies:

7 "Deskill" is defined as "to reduce the level of skill needed for a job" in Merriam-Webster Dictionary (2024). Deskilling is a common term used in the context of automation to refer to this effect related to the introduction of new technologies into a job. For seminal discussions of the concept, see social scientist Harry Braverman's (1974) chapters "Machinery" and "A Final Note on Skill" in his book *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century* as well as James Bright's (1958) chapter "The Impact of Automation on the Work Force" in his book *Automation and Management*.

> *At the end of the day I'm a tradesperson, like, I'm gonna want to keep my trade going, you know. If we start having robot[s] do the trade, then you're not gonna need as many plumbers, so that's gonna be other people working… (Plumber 4)*

Fitness of robots and sites: Based on prior experience with different technological equipment on job sites, some workers shared reservations about the future adoption of robotic systems on sites. These reservations revolve around two critical questions. The first is whether these technologies are ready for construction sites in terms of their reliability and safety. The workers still have a degree of mistrust. "There is always mechanical equipment breaking down, so that is my only concern [about robots]," said Cleaning-Housekeeping Worker 1, who is also an experienced equipment operator. "But if that would fail, it would kill somebody" was Sheet Metal Worker 1's worry, voiced in a similar skeptical tone, about a robotic duct jack we discussed in a hypothetical scenario. Plumber 4 raised an essential point. Comparing robots to other "machines that get beaten" in construction, he asked how much abuse these "highly computerized" machines could take, and who would be "servicing them."

The second question in the workers' reservations concerned whether construction sites are ready for robotic systems to operate smoothly. One thing the workers know well is that while their labor changes the site every day, the changing site conditions the way they perform their labor. This is evident in the example of frames and material stacks constraining the worker, lift, and material circulation on Level B. These same constraining elements would also condition a robot-based operation unless things on sites are rearranged for those operations. Sheet Metal Worker 1 pointed out the need for this type of rearrangement based on the example of a commercial-grade drilling robot.[8] Imagining this machine in the middle of the material stacks on Level B, he stated: "that [robotic drilling] could work, if you do it before all the materials are installed." Along the same lines, Framer 1 brought up the many variables on sites that might make using a robot complicated:

> *There's too many variations I think that go into it to use a robot… You never know what's coming up. I mean, you can't control the rain, even a big puddle on the floor, or a sub [subcontractor] being in your way, or the engineer making a wrong mark, and you have to stop what you're doing, go to your boss, and he has to put an RFI [request for information] out. (Framer 1)*

In his description, a construction site is presented as not only a shared workspace but also a large-scale "coordinative artifact"—borrowing from sociologist Kjeld Schmidt and computer scientist Ina Wagner (2004)—combining information, materials, and physical conditions that currently shape the work and have to be considered carefully for robot integration.[9]

Roles for robots: Lastly, the workers' general views about the adoption of robots were mixed, rarely black or white, and mostly nuanced. Their views lay on a spectrum between openness on one end and skepticism on the other. For instance, Sheet Metal Worker 1 was on the skeptical side. He thought that these technologies would bring unnecessary complications into his task, asserting: "having battery tools is enough, that's all the technology I want." Some other workers gave robots the benefit of the doubt. Ironworker 2 was less skeptical than Sheet Metal Worker 1 but acknowledged possible complications. He knew that "the setup for some of these complex systems,"

8 See "Jaibot Drilling Robot" by Hilti Corporation (2020).

9 With the term "coordinative artifact," Kjeld Schmidt and Ina Wagner (2004) describe various instruments that link and provide order to individuals and collaborative practices, based on their ethnographic study of an architectural office. Aside from the 2004 article titled "Ordering Systems: Coordinative Practices and Artifacts in Architectural Design and Planning" that I primarily reference, Schmidt and Wagner (2002, 2003) also used the term in earlier publications of the study.

as opposed to “just hiring a [crew], doing the old way,” is sometimes a mental barrier against their integration. Yet, he also argued that a transition should not happen overnight. According to him, there is no need to “jump right out of the gate,” and “that’s why we have version 2.0 with everything, because things get better with time.” With the same open mindset, Plumber 4 suggested that robots could be useful if the right role for these technologies were found. He supported the idea of “things assisting” and believes that such technologies may make his life easier. Still, he was against the idea of technologies “taking away the trade.” He worried about the idea of robots as substitutes. Similarly, Cleaning-Housekeeping Worker 1 was open to the use of robots in less dexterous and laborious activities, which fits what these technologies are designed to do: “if it [a robot] can come in and clean as good as I can, I let it do it.”

In sum, the ethnographic study of the New Scaife Hall construction provided an understanding of how work is organized and enacted on a floor. The key insights we gained as a group were the following. Based on task location and coordination, trades placed their materials in designated zones, which restrained movement. Workers (in units of two to three) performed tasks in planned work zones, and remained in their zones as best as they could to minimize safety risks and conflicts. Interactions within each of those units were more common than they were between the different units of a trade or units from different trades. Further, the ethnographic study detailed the qualities of the observed tasks, grounding the high-level, generic descriptions of construction work based on three insights. First, one sees certain qualities only when one watches the enactment of tasks or listens to workers. For example, the arrangement of tasks is seemingly messy but not unorderly, as I showed with the coordination of work zones. Similarly, next to physical labor, tasks involve many hidden intellectual layers. Second, observing tasks allows us to see that not every one of them possesses the general qualities associated with construction work to the same extent. Third, contingent on the role of workers and materials, different degrees and paces of mobility along with different degrees of communication occur across them. Finally, the ethnographic study captured the workers’ mixed views about robot adoption in construction, redirecting our attention to the situated nature of tasks, skills and deskilling, fitness of robots and sites, and the appropriate roles for robots in construction as key issues to consider in technology design.

Chapter 4

Reimagining the Role of Robots: The ReAC Project

4.1 Outline

Following on from the previous chapter, I draw two parallel threads from my participatory observations and analysis of ReAC's technology design efforts in this chapter. In one thread, I discuss the ways the New Scaife Hall study worked to combine robotics and construction knowledge in ReAC's technology design work. I show this by describing the three roles that the specific elements of the study played in conceiving and developing a system. The first one is that the task descriptions and the data viewer brought immersion, imagination, and conceptual guides next to technical knowledge into conceiving a robot. The second one is that the knowledge of zones and interactions, the task descriptions, and the worker perspectives guided us to focus on the many activities that constitute tasks, consider their specific qualities and ways they are enacted—rather than their general character or commonalities—, and consequently evaluate a rich set of possibilities for mobilizing RL towards worker support. Seeing and analyzing work in this way allowed us to notice ancillary activities as opportunities of robot support in tasks, and situate a mobile robot prototype with a support role, tool-hardware delivery, in a specific carpentry task. The last role is that the knowledge of zones and interactions, specifically the specific site conditions, obstacles, and density of people, informed the development of this prototype.

In the parallel thread, I look into the group's conception of construction work and automation that underlie our work. I identify that, drawing from an understanding of how work is enacted on site and a sense of how workers view robots, we conceptualized construction work as richly layered with parts that workers invest time, use skills, or attach value. Based on this conception, we conceived and demonstrated an alternative approach seeking to support workers' time-invested, skill-invested, and valued activities by automating their ancillary ones, which I view as alternative to the historically prevalent task automation ambitions. In conversation with the entire team, this approach helped shape the RL-based tool and hardware delivery robot to support carpentry formwork, which displays an alternative support function next to archetypal single-task robots in construction—that are designed to perform primary skilled activities. I discuss this as the group's reimagination of robots' role in construction.

4.2. An Ethnographic Study of Robotic System Development

4.2.1 Mill 19

At some point on the Penn-Lincoln Parkway driving towards Oakland, a large steel frame with an army of solar panels on top appears in the distance. Standing alone on a brownfield in Hazelwood along the Monongahela River, this building, called Mill 19, holds the traces of Pittsburgh's two intertwined technological histories. The steel frame is an artifact of the city's heavy industrial past, left over from 1943 when the rise of its developer and first owner John & Laughlin Steel Company was nearing its end (Almono Limited Partnership, 2022).[1] Since the steel operations stopped in Hazelwood in the late 1990s, the brownfield site and the frame have changed hands and eventually become part of an ambitious urban redevelopment effort (Almono Limited Partnership, 2022) (Figure 1). Through this effort, Mill 19's steel frame and CMU's tech-entrepreneurial ambitions since the early days of the Robotics Institute, crossed paths. Driving closer, one can see that the frame shelters a modern, box-shaped research facility, which hosts CMU's Manufacturing Futures Institute (MFI), a manufacturing research catalyst, and the Advanced Robotics for Manufacturing, a non-profit

1 These details are from a brief titled "The site's history," published online by Almono Limited Partnership (2022), the landowner of Hazelwood Green site and the visionary for its redevelopment including the Mill 19 building. See the source in Bibliography.

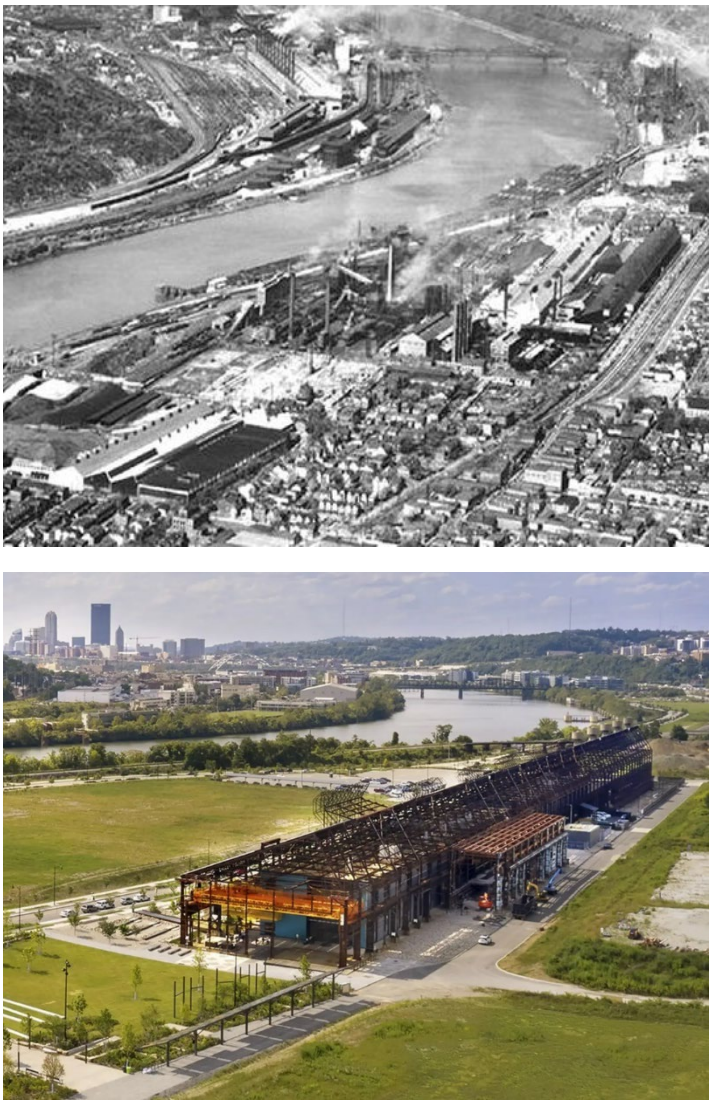


Figure 1. Hazelwood in the early 1900s (top) and post 2019 (bottom) (Sources: Almono Limited Partnership, Carnegie Mellon University)

research partner of the university, along with various faculty-led research groups (Figure 2).[2] Aside from the offices of these two major research initiatives, the box also contains an additive manufacturing lab, a machine shop, flexible lab spaces, meeting rooms, and several lounge-like workspaces (MFI, 2022).

"Carnegie Mellon University researchers and robots will soon occupy a state-of-the-art facility at the newly constructed Mill 19 building," the university announced the facility's opening in the fall of 2019 (Carnegie Mellon University, 2019). We, as the ReAC group, joined these researchers in 2021, when MFI granted us funding and a lab space in a portion of the building called the "Low Bay Project Space" (Figure 3). This lab space, along with a few meeting rooms, accommodated ReAC's weekly meetings, biweekly site information sessions, and a major part of the technical development work informed by the New Scaife Hall study. In my role leading the ethnographic study, I actively

2 For more information on CMU's Manufacturing Futures Institute (MFI) and the Advanced Robotics for Manufacturing (ARM), see www.engineering.cmu.edu/mfi/ and www.arminstitute.org.



Figure 2. The Mill 19 facility (Source: Carnegie Mellon University)

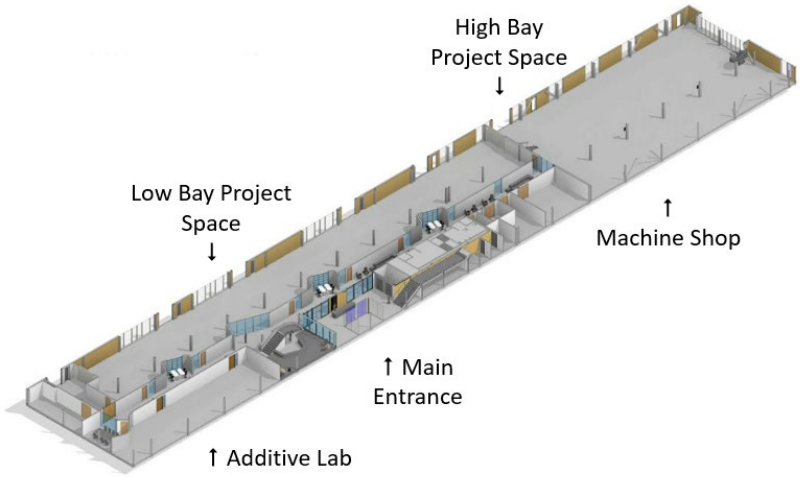


Figure 3. The Low Bay Project Space in Mill 19 (Source: MFI)

participated in these various activities. Yet, Mill 19 meant more than just a workspace for me; it was also an ethnographic site (in addition to the New Scaife Hall construction) where I closely studied the ReAC's system development activities and investigated our conceptions of construction work and automation in construction within these efforts.

4.2.2 Ethnography Drawing Robotics and Construction Closer

Once almost every week in 2022, the ReAC group would be on a large screen in a meeting lounge or in a classroom at Mill 19, carefully following a presentation of a building task from the New Scaife Hall construction in a site information session or a presentation of the latest system development updates in a technical update session. The site information sessions aimed at familiarizing the group with the tasks that I observed on the construction site and the technical update sessions were geared to acquaint them with the technical explorations conducted in the lab. These sessions were

also for creating an exchange between the ethnographic and the technical sides of the project. They featured rich interdisciplinary technology design discussions driving the system development efforts.

Based on my participatory observations and analysis of these discussions and development efforts, I discuss two parallel threads in the next three sections. In the first one, I unpack the ways ethnography drew robotics and construction closer in system development work. I do this by foregrounding the three roles that the specific elements of the New Scaife Hall study, the task descriptions accompanied by the data viewer, the knowledge of zones and interactions, and the workers' perspectives, played in system development. In the parallel thread, I identify our conceptions of construction work and automation in construction that underlie our work as they were shaped by these specific elements.

Bringing the Site Into the Robotics Lab: Prototype A – Data Viewer

Each day of the New Scaife Hall study provided a rich, diverse raw data set coming from different sources, in different modalities, and in great quantity, amounting to pages of notes, three time-lapse videos, up to one hundred photos, a couple of plan drawings, and numbers indicating work interactions.[3] How to organize, review, and communicate a sorted version of this set with the group almost as I collected it and how to navigate this set effectively and interactively were the reflexive questions that guided Prototype A, the data viewer I built as a response. Because I could not bring the members to the New Scaife Hall site, I tried to bring the site to them through the viewer.

The prototype is a combined file management system and interactive dashboard developed in Processing, which combines a Java-based graphics library, an application programming interface, and an integrated development environment.[4] It imports the data items from the local directory, sorts them into "rooms" representing the days of the fieldwork, and allows for interactively navigating these rooms to view them (Figure 4). Accessible via buttons, the rooms include "video pages," each containing a player showing the recordings of tasks focused on that day, a key floor plan delineating the camera location and the zone where the recorded task was performed, and a workflow diagram summarizing the key activities performed by workers (Figure 4). Linked to these video pages are "photo pages" containing a custom viewer displaying photographs of the key moments of the task shown as well as the tools and materials workers used (Figure 4). Each room also features a "site page" including a diagram demarcating the work and material zones on the floor, and some rooms display a network graph visualizing work interactions within and between the building trades recorded on that day (Figure 4).

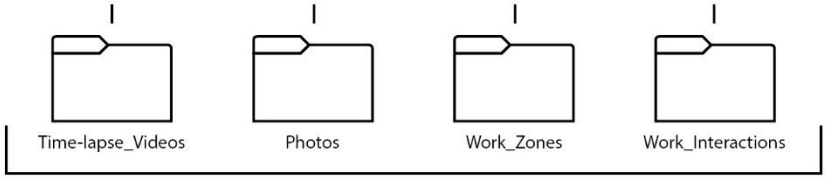
On Site With Multiple Task in Mind: Collective Imagination and Conceptual Guides

Compared to most robotics projects, the ReAC group worked and thought not only with the state-of-the-art methods and models but also imagination and conceptual guides that stemmed from our initial premise to "rethink automation" and conversations during the site information sessions. In these sessions, I presented the descriptions of tasks accompanied by scenes on the data viewer

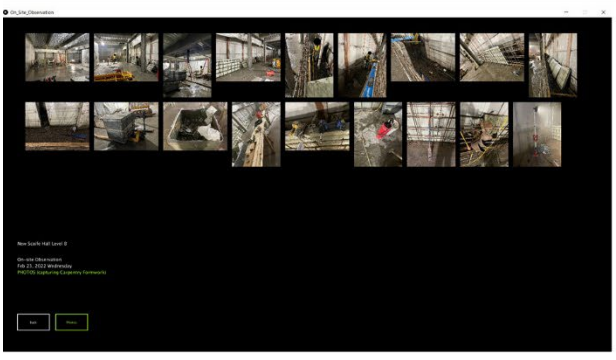
3 This raw data set was sorted and stored in folders by date and subfolders by data type within a typical local directory. Even after sorting by type, the size of some data, such as photos, was still too large to digest, or the quality of some data, such as work and material zone sketches, was naturally low and not easily legible for viewers. Therefore, a round of manual data processing, which involved systematically selecting twenty to thirty photos that could help describe a building task, digitalizing the work and material zone sketches, and formatting the work interaction data, turned out to be an essential step in preparing these items for review. These sorted items could then be accessed by navigating the directory and reviewed via multiple application interfaces, such as the image viewer or video player pre-installed in the operating system. However, navigating and presenting data through these multiple interfaces (as opposed to a single interface) is often not practical.

4 See www.processing.org for more information on Processing (Fry & Reas, 2001).









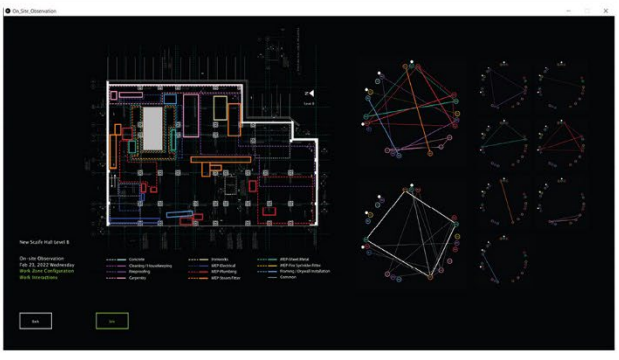


Figure 4. The data viewer: an example video page (middle), photo page (bottom left), and site page (bottom right)

showing how they are performed. These descriptions and scenes facilitated an engagement of the ReAC group with construction work, which was distinctively characterized by seeing and immersion. Specifically, listening to the detailed descriptions and seeing helped immerse the members in the tasks and the site. Discussions supported by these verbal and visual materials involved a collective imagination in conceptualizing systems, which I noticed in the sessions. They began to bring the members' growing knowledge of tasks and existing knowledge of robotics along with their conceptual and technical angles together. Additionally, these discussions gave rise to a set of themes, which conceptually guided our conception of systems. Alongside the technical artifacts, all these became our technology design instruments.

The descriptions and the data viewer providing the original scene allowed us to picture imaginary systems as part of the workflows on the screen. For instance, after learning about and seeing the fireproofing task on the viewer, Researcher 2, who led the technical side of the project, imagined a concept navigation system, which she then used to build a line of thinking. She began by delineating the system as part of the fireproofing workflow, performing concurrent cleaning, an activity done by the supporting worker I covered in the previous chapter. She described scenes from a potential workflow using this system, in which she began to envision its capabilities: "the robot will be first trying to avoid the area where [the spray] debris will be falling… and then it will be cleaning." The system, in her view, could adaptively position itself relative to the lead worker on the lift to avoid the dripping spray. Following this description, she switched to a more technical mindset and pointed out the complexity of this idea by translating this imagined capability into a perception problem: "either part [cleaning or spraying] will involve that the robot understands the whole height, but that can be quite difficult in terms of perception, in terms of how [the] robot perceives the height."

Various exchanges in sessions showed how we weaved the knowledge of tasks (coming from the ethnographic study) and of robotics as well as conceptual and technical lenses together. Combining two knowledge worlds and angles enabled us to imagine technicalities of systems in conjunction with the reality of tasks and evaluate. For instance, based on her grasp of the task scene in front, Researcher 2 expressed concern over two technical challenges regarding the system she came up with: how the system would detect the lead worker up high and, accordingly, circumscribe the extent of navigation and cleaning. What drove this concern and foresight was a technical angle she often injected into these conversations to remind us of practicality. At this point, Researcher 1, who led the project, responded to Researcher 2's concern by building on her idea from a more designer-oriented angle taking into account the physical elements on site. He suggested considering "ways for the human [worker] to demarcate and delineate what the robot will do" and markers "as a way of leaving instructions for the robot [on the floor[" which were the type of ideas we later described with the term "assisted automation" in conversations (Cardoso Llach et al., 2017). The example he gave was the lead worker "spray paint[ing] fluorescent green on the start and end points of the beam" to direct a manipulator to spray.

Bringing together these two types of knowledge and angles at times created a mild and, in a way, productive tension between conception and technical requirements.[5] An example instance was during the fireproofing discussion. "So, the robot will be operating on that height?" was Researcher 2's response to Researcher 1, who suggested a system to support the spraying work on the lift as an alternative to the cleaning system we discussed a few moments before. "A bit scary" was Researcher 2's initial response, packed with safety concerns, because the warehouse robot the

5 This tension was due to differences in intellectual inclinations, viewpoints, and technical expertise as well as design and computing literacy among the group members. It surfaced more in certain activities, such as the site sessions, core group meetings between Researchers 1 and 2 and me, and conversations between Researcher 2 and myself as the research leads.

group planned to use at that time—a Fetch Mobile Manipulator—was a complete mismatch for the task. In addition to her point, I brought up possible sensing and manipulation challenges for a robot to handle intricate spraying movements in tight angles. Conversely, according to Researcher 1, if we went for this concept, we could opt for an appropriate chassis and program a manipulator to spray easily accessible and larger surfaces marked by the worker. “We could potentially leave the question of height to a later stage,” he said, signaling that it would be ok to initially demonstrate the concept and the system on a low-height setup. This approach was to stay conceptual at this early stage and not to let foreseen challenges in implementation limit our view of the site and ideation too soon.

This balancing between imaginary systems and the discrimination of them often took place between Researcher 2 and myself as well, since she would immediately ground the potentials I found in tasks in a technical viewpoint. One of those potentials I proposed to her concerned nailing bottom tracks in framing. To Researcher 2, developing a system for nailing would involve the challenge of handling navigation and manipulation together in one project. Further, any scenario where a manipulator shoots towards a target would make precision a hard problem to solve and bring safety risks. In the same vein, I would evaluate Researcher 2’s ideas, which nicely lent themselves to the RL exploration she was inclined to but did not always offer practical value for construction. For instance, we moved away from her idea of a system that delivers tools and hardware from a central repository to all trades and between trades based on a key insight I described in the previous chapter. In practice, trades keep their tools, hardware, and materials exclusively in their zones and limit their interactions. Thus, there was no demand for such central distribution or sharing on the site. As seen in these examples, technical and ethnographic knowledge stood antagonistically at times and provided balance.

In addition to immersion and imagination that allowed the members to weave robotics and construction, the New Scaife Hall study led to the emergence of a set of themes that worked as conceptual guides as we ideated systems. These were the notion of robot accompaniment that helped us to identify our positioning of robots and a set of values, namely task specificity, credibility, and technical feasibility, which we sought to have in our ideas.

Task descriptions, in particular, the notion of supporting roles, that I described to my colleagues inspired Researcher 2 to associate those roles with the notion of robot accompaniment. This notion came up in the fireproofing discussion and remained a guiding theme in conceiving a system. Fireproofing was the first task involving the typical ground and up roles that I presented to the group. This sharing of roles between Fireproofing Workers 1 and 2 that I observed on Level B caught Researcher 2’s attention too. “What is kind of attracting me is that there’s this company, so there are basically two people working together,” she said, and, building on that point, she went on to frame an analogous relationship between a worker and a robot. “The accompany kind of thing,” as she described in the moment, or “robot accompaniment,” the robotics term she used later, nicely conveyed the role she began to imagine for a robot. Building on this notion, she then imagined a scenario in which a robot would “follow [Fireproofing Worker 2] around,” offering some form of needed “assistance” while the worker is “pushing the debris or doing any cleaning.” This “basic functionality,” as she called it, could be a way of demonstrating the concept of a robot accompanying a worker. “Or vice versa… the person accompanies a robot,” Researcher 1 chimed in, prompting us to think the other way around. He reiterated the assisted automation scenario where robotic cleaning or spraying is guided by worker demarcations. In this form of accompaniment, the worker would be “minding” and “fine-tuning” the robot’s “execution,” a role similar to that of Fireproofing Worker 2 while supporting Fireproofing Worker 1, as he highlighted. Accompaniment as a frame helped us crystallize our positioning of a robot in construction: a robot could assist or be assisted.

The knowledge of tasks' specific qualities that I presented in detail to my colleagues allowed us to evaluate our ideas on the basis of specificity, credibility, and technical feasibility. These themes became the guardrails to keep the concept systems remained aligned with our goals. An important goal was to develop a contextual system in the sense that it would cater to a need in a specific task in a credible way. Therefore, task specificity and credibility were two values we sought in each idea. Researcher 2 and I sometimes were inclined to think alternatively about task specificity. For a while, we considered whether one form of accompaniment could work for various tasks. It was Researcher 2 who first shared this notion with the group when she brought up the term "accompaniment" during the fireproofing discussion. "I'm thinking in more like [a] general way," she said, following with an intriguing question: "so what if it's not fireproofing?" To her, a technical exploration could be built around one basic functionality that would demonstrate accompaniment, but instead of one specific task, that same functionality could perhaps be applicable to multiple tasks. Researcher 1's response to the idea was gentle pushback: "we probably want to strive for a scenario of application that is not too general," he affirmed. According to him, "accompaniment" would be a strong concept "as long as it is connected to a specific workflow."

In a part of the same discussion focusing on rebar tying, I revisited this idea. Carrying and delivering ironworkers' tools—as opposed to workers carrying tool harnesses—was a preliminary use case that I pointed out to Researcher 2 before this session. I proposed that this kind of support could be transferable to some other tasks: "it could be for rebar tying, in another case, maybe it can be along this [same] wall for the formwork guys." I emphasized that a robot supporting multiple tasks with one functionality and without building anything would be an alternative approach to the common vision of robots as surrogate builders, tying rebars, picking and placing bricks, or installing panels. Researcher 1 sought further credibility in the specificities of these tasks. He reasserted that the "relevance [of any system we propose] to the specifics of [the] construction site" should be "beyond their generic application to any of the tasks." "If it [accompanying] is about carrying particular items from one place to another, why those items," and "why that… in this specific context becomes essential" were the key to building a case for the idea. As these two examples that I described shows, seeking task specificity and credibility kept our ideation in line.

The other important goal was to be able to develop a functional system that we could demonstrate in the existing conditions of a site. Therefore, technical feasibility of a concept system, in addition to its task specificity and credibility, was another important value we sought. This was Researcher 2's top concern, as expected. While she fully acknowledged the significance of "a credible scenario," whether "we have the technical power" to make it work within the timeline was the critical question she first raised to Researcher 1 and me in a core group meeting. Because her inclinations were towards navigation-related work, this is where she felt she had enough technical expertise and could help other researchers develop. For his part, Researcher 1 was still open to ideas that also included manipulation at that early stage. Although we had not decided on a robot or any other hardware to work with, he was against the concerns of technical power and time overriding our imagination. "What matters is a credible scenario," he said in the same meeting to re-emphasize our priority. "It can be a wilder scenario that can be carried into a longer timeline and agenda," he continued, adapting his tone to ease Researcher 2's concerns and help fuel her imagination. In such moments, balancing credibility with feasibility became the other guardrail during ideation.

A Brief Description of System Development Work During Ethnography

Over the course of two years, a list of members from the interstices of architecture, computational design, HCI, machine learning, robotics, and mechanical engineering contributed to the ReAC

project (Figure 5). On the technical side, the project was biased towards exploring the potential of RL in construction robotics. Whether the focus would be manipulation, navigation, or both was not fully determined at first. Integrating manipulation, “a little beyond navigation,” in Researcher 1’s words, was still a possible consideration as the New Scaife Hall study was ongoing. We even discussed a couple of combined systems, including Researcher 1’s suggested fireproofing system detecting cues from workers which I previously noted in this chapter, and Researchers 4 and 2’s rebar tying system positioning the rebar for workers to tie. Still, for Researcher 2, navigation was the priority. Manipulation should only come into the picture if there was an appropriate scenario, time, and resources. With that, she initiated the hands-on efforts by configuring a navigation system. Running parallel to the New Scaife Hall study, this work included exploring different robot chassis, designing a system architecture, and selecting sensor hardware for the system’s perception modules.

Researcher 1	Principal investigator —qualitative research, computing+computational design, and experience in the AEC industry	Researcher 7	Master’s student —computing+machine learning, robotic fabrication, and previous training in architecture
Researcher 2	Technical research lead (Ph.D. candidate) —computing+machine learning, robotics, and experience in the AEC industry	Researcher 8	Master’s student —computing+machine learning and previous training in architecture
Researcher 3	Master’s student —computing+machine learning, robotics, and previous training in architecture	Researcher 9	Master’s student —computing+machine learning and previous training in architecture
Researcher 4	Master’s student —computing+machine learning and previous training in architecture	Researcher 10	Master’s student —computing+machine learning and previous training in architecture
Researcher 5	Undergraduate student —architecture, computing+computational design, and robotic fabrication	Researcher 11	Master’s student —computing+computational design and experience in the AEC industry
Researcher 6	Master’s student —mechanical engineering, computing+machine learning, and robotics	Myself	Qualitative research lead (Ph.D. candidate) —qualitative research, computing+computational design, HCI, and experience in the AEC industry
		Researcher 12	Co-principal investigator —computing+machine learning and robotics

Figure 5. Members and their knowledge they brought to the ReAC project

With all these parts tied to each other, this work was time-consuming by nature and the challenges we faced made it take longer. For instance, due to technical issues, we had to switch between multiple chassis during the ethnographic study. This delayed sensor purchases and experiments, therefore, the team mainly focused on the system architecture. The initial architecture was designed by Researchers 2 and 6 soon after the study kicked off. This architecture helped Researcher 2 arrange the division of labor. Researchers 3 and 4 took on the configuration of the system’s localization and mapping modules. Researchers 7, 8, 9, and 10 started with the experiments to find the best method and model for the people detection and tracking module. Researchers 2 and 6 worked on the experiments to select global and local path-planning methods for the navigation modules. Researchers 2 and 6 also carried out the RL module. All these members had much to cover in terms of new technical knowledge, and, consequently, every module took significant time and labor. Based on the architecture, a sensor suite for navigation including LiDAR scanners, RGB cameras, and an inertia measurement unit (IMU) was obtained. Nonetheless, their integration had to wait until the issue of the chassis was settled. As more and more time was spent on the challenges

in configuration, the group had to strategically drop the possibility of a manipulation system and concentrate on a navigation-only system instead.

On the other hand, configuration taking longer than planned provided certain advantages. For instance, the architecture that Researchers 2 and 6 put together clarified the capabilities our system would have. We all began to think on this basis while conceiving concept systems. Next to our conceptual guides that I described in the previous section, the architecture became a sort of technical guide in this effort. Meanwhile, not having a dedicated chassis for a period enabled us to think of the system's appearance freely. What a potential system might look like to perform a function we imagined and how it might withstand the site conditions became important considerations. Perhaps more importantly, I exposed the members to the knowledge of the site, tasks, workers I gathered before they dived deep into implementation during this period. Thus, delays in configuration allowed us to be more contextual.

Robots for Ancillary Activities: Practical Guides, Use Cases, and a Scenario

For the fifteen weeks that I was active on the New Scaife Hall site, we worked with a few valuable sources—the knowledge of zones and interactions, the task descriptions, and the system architecture—to find use cases and conceive concept systems. The latter half of this period was when I also interviewed workers, which brought another valuable source, their perspectives on robots and feedback on some of our specific ideas. Guided by these sources, we found and evaluated a list of robot support use cases. We eventually narrowed down to the tool and hardware delivery use case as technically feasible for the group and credibly useful to carpentry workers. We situated and virtually tested this idea within carpentry formwork task through a scenario, and identified the robot navigation capabilities needed to realize it.

By looking closely into these efforts, in this section, I expand on the second role that the specific elements of the New Scaife Hall study played in system development. I identify how the knowledge of zones and interactions, the task descriptions, and the worker perspectives I described in the previous chapter guided us in our analysis of construction work: staying at the level of activities, attending to their specific qualities and specific ways they are enacted—rather than their general character or commonalities—, and exploring a rich set of possibilities for mobilizing RL towards worker support. I show that this close view of tasks allowed us to notice ancillary activities, such as garbage delivery, cleaning debris, rebar positioning, or tool-hardware delivery, which are only visible on site, and shift our attention to these while searching for use cases. Relatedly, these ethnographic materials guided us to conceive and demonstrate a support role for a robot, handling tool and hardware delivery in carpentry formwork task.

Drawing from these, I discuss our conception of construction work as richly layered with activities for which workers invest training time and skills, and intellectual parts to which they attach value. Related to this conception, I situate how we built a view of automation seeking to mobilize RL to automate ancillary activities and support workers in performing their time-invested, skill-invested, and valued activities. I consider this alternative to the historically prevalent task automation ambitions in construction. Lastly, I position the support role we conceived and demonstrated with the final prototype, alternative to the lead role appointed to archetypal single-task robots, as a reimagination of robots' role in construction.

The key insights of the New Scaife Hall study, specifically zones, interactions, and worker perspectives worked as practical guides in searching for and evaluating use cases. For instance,

with potential issues of occlusion, a large variety of obstacles to detect, and the lack of navigation paths, we concluded that without reorganizing material zones, operating a robot across the whole floor would be unworkable. Similarly, workers' movements confined within their work zones to maximize safety and minimize conflicts, and their interactions mostly taking place in a typical trade unit of two to three workers in a particular zone made us look for a case within these units. Because floor-wide navigation is impractical, worker movement between zones is limited, and worker exchange across zones is rare, limiting robot support to one specific work zone and a typical unit of two to three workers of a specific trade made sense to us. Therefore, every use case we found came from a task carried out by a unit of a specific trade, namely rebar tying, MEP ceiling installations, drywall installation, fireproofing, cleaning-housekeeping, framing, and carpentry formwork.

On top of how we imagined robots might be useful, the workers' perspectives allowed us to sense where they perceived these technologies as useful or otherwise. As I described in the previous chapter, an important point the workers made concerned the situated nature of their tasks. They had the preconception that, like on an assembly line, integrating a robot would typically impose a new scripted order onto their tasks. Part of the idea behind our project was to resist this by conceiving and developing adaptive alternatives fitting into workers' natural course of action rather than completely reconfiguring it. We discarded a few use cases (drilling, rod installation, and material lifting) we considered for MEP ceiling installation based on our consideration of this point and feasibility of developing the precision and special hardware potential systems would require.[6] Systems we conceived for these activities would bring in a new workflow and working with these machines would require special planning and a reconfiguration of work zone. The workers also had a related skepticism about the fitness of robots for sites. Without any rearrangement of sites, the typical congestion on the ground and up high could make the operation of these machines unsafe. We eliminated the layout use case that Researcher 2 and I considered partially for a similar reason. A layout system would require clearance and reconfiguration on sites.[7]

The workers' other significant point is related to skills and deskilling, which is also tied to the role they deemed more appropriate for robots. The workers viewed some activities, in which repetitive-mechanical, skilled, and intellectual parts are strongly intertwined, as unsuitable for a robotic system to perform. An example is determining configurations, drafting, and cutting drywall panels. Although the activity is seemingly mechanical, the way a configuration is determined is based on "just experience" and "talking back and forth [with each other]," in Drywall Worker 1's words.[8] The improvisational nature of it that I observed on site, together with Drywall Finisher 1's confirmation of this quality, made me pass on this one quickly. The size and weight of panels and the dexterity required to handle them, which I observed on Level B, led me to a similar judgment about

6 For instance, robotic drilling would bring the challenge of combining manipulation and navigation to achieve high precision. Additionally, an advanced semi-autonomous drilling system already existed in the industry. See "Jaibot Drilling Robot" by Hilti Corporation (2020) also mentioned in Chapter 3. For rod installation, the challenge was fastening—in addition to precision—, which would have to be handled as a complex robot manipulation problem or by developing a dedicated mechanical system. Similarly, line fitting involves working with tall and thin materials such as electrical conduits, tall and thick materials such as steam pipes, and large-scale materials such as ducts, more often than not in tight, high spaces. A form of support would be to robotically lift these materials, which would again involve developing a designated mechanical system to handle these items.

7 Seeing Plumber 1 mark the hanger points on the floor made me think of this system that could follow him along the lines on the blueprint and draw. Later, Researcher 2 came up with a slightly different version, which would directly laser-mark the hanger points on the ceiling instead of marking them on the floor. She knew that "laser pointing" may demand higher accuracy in robot localization compared to "following a person" and was skeptical because this "could be a challenge." Other than these technical challenges, I was hesitant, because such a system already existed as a commercial-grade product—see "FieldPrinter" launched in 2019 on Dusty Robotics' (2023) website.

8 Drywall workers take decisions on the spot to ideally accomplish two goals. One is to maximize the size of each cut for overall efficiency: "ideally… I wanna give whoever is the up guy the biggest sheet possible," Drywall Worker 1 explained. The second one is to pass drywall finishers, who come after installers, the minimum number of joints in each section: "I could cut it [a panel] into ten pieces, but then that's ten joints that they [finishers] have to mud… so ideally, I want to give two pieces every time."

installation. Positioning panels is a fairly complex manipulation problem, which we only see in contrived demonstrations.[9] Instances of this activity on an actual site involve the opposite: moving around obstacles, positioning custom-cut panels, and working in high sections constricted by MEP lines during topping out are all part of the job (Figure 6). This makes robotic drywall installation a complex, and perhaps impractical, project. Even if certain repetitive-mechanical aspects may be suitable, some workers still want to handle them themselves because they value their trade. For instance, although machines for robotic drilling or rod installation would be helpful to take a considerable load off of workers, these machines capabilities are replicative than supportive, which some workers believe would cause deskilling.

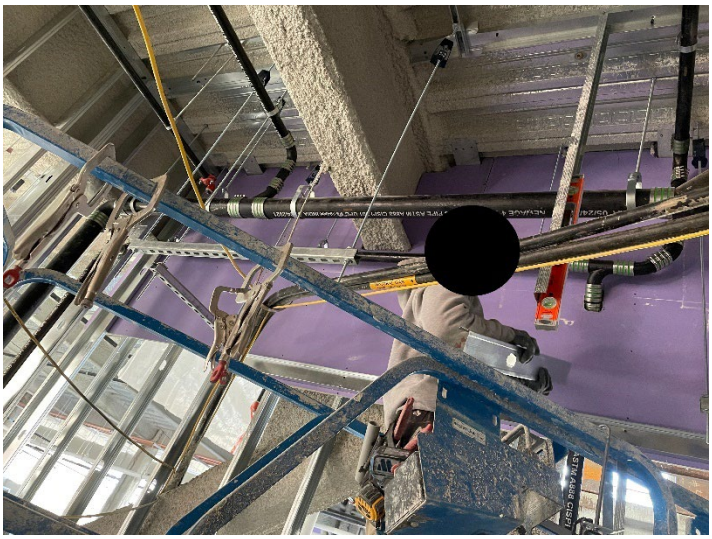


Figure 6. Drywall Worker 2 performing installation on a lift in a constricted area

Workers also provided direct feedback on some of our use cases. For instance, tool-hardware delivery with a navigator following the ironworkers tying rebars seemed to be a supportive capability. While such a system could eliminate workers' tool harnesses, I debated whether it would be overkill for a robot to carry only the pliers, wires, and knife they use or whether it would slow them down. "The tools that we have are on us [to keep up the pace]," Ironworker 2 stated, resolving the debate (Figure 7). Robot tool and hardware delivery was also considered for keeping tools and hardware in proximity to MEP worker pairs during ceiling installation. MEP workers saw lifts as more convenient and safer when handling large-size materials: "usually, the connection [up] guy is on a lift, because he is more stationary bolting everything together," said Sheet Metal Worker 1 (Figure 8). "You can have a little tray up there with your tools," said Plumber 4. This detail implied that a robot moving stuff would be redundant next to a lift. "Sometimes, they [lifts] don't fit," and sometimes "you can't go up because you got rods hitting it [the lift]," said Plumber 4 and Fire Sprinkler Fitter 1, respectively, and that is also when ladders come in handy (Figure 8). The tradeoff is that workers cannot keep their tools up on ladders, and that is where a tool-hardware delivery system could help. However, because ladders are preferred in tight spaces, the paths for the system would be limited or fully obstructed. The same idea was also shared with the drywall installation workers. "When you're doing bottoms, you just want as little on you as possible," Drywall Worker 1 stated (Figure 9). These discussions revealed that these use cases might not be viable.

9 For instance, see the demonstration of "HRP-5P," a humanoid robot that can lift and position panels developed by the National Institute of Advanced Science and Technology (AIST) of Japan (2018), in a contrived set up.



Figure 7. The tool harness that the ironworkers used to keep tools at reach

Other than the workers' perspectives, the granular view of tasks that the New Scaife Hall study brought to the group's attention helped in couple ways in understanding construction work and conceiving systems. As I discussed in the previous chapter, every task or activity that constitutes a task embodies the qualities that we know about construction work to a different extent. To reiterate an example, spraying in fireproofing is more skill invested, dexterous, and reliant on routine mechanical moves than the cleaning portion of the same task, which is unscripted and depends less on dexterity and more on intuitive physical cleaning moves. This view of tasks helped us to start analyzing construction work at the level of activities rather than tasks as we searched use cases or conceived systems. It guided us to look for not just the primary but various activities in each task and consider them based on their many qualities—rather than the general character of tasks, commonalities between tasks, or merely the basic repetitive activities in tasks.

This signals the group's conception of construction work as a layered practice with many activities, specificities, and various qualities to different extents. We recognized that this layered constitution consists of activities for which workers invest training time and skills, intellectual aspects to which they attach value, and parts that are suited to machines. The way of distinguishing activities (based on their qualitative character) in the instances that I offer next is underlined by this conception. In these instances, we recognized activities that a machine could perform to support workers.

For instance, we did not only pay attention to the primary activity of spraying in fireproofing. The other layer of the task, cleaning, was also seen as worthy of discussion for robotization. With its less dexterous and more physical nature, this activity was immediately seen suitable for a robot, as Researcher 2 and I remarked at different times. We both imagined a navigation system supporting Fireproofing Worker 1 by cleaning or Fireproofing Worker 2 during cleaning. On the other side, certain parts of spraying were also considered as appropriate for robots to perform. When I broke down Fireproofing Worker 1's routine into "straightforward" left-and-right moves to spray the wide flat parts of beams and "intricate" moves to spray the bottoms and corners, Researcher 1 had a grasp of the dual quality of this work. Accordingly, he brainstormed a complementary role share between a robot handling the straightforward axial move to spray "the flat parts" and the worker handling the intricate moves to spray "the corners." Distinguishing activities and qualities was also at play during the rebar tying discussion. Not far from my intuition, the group viewed wire tying as a highly complex

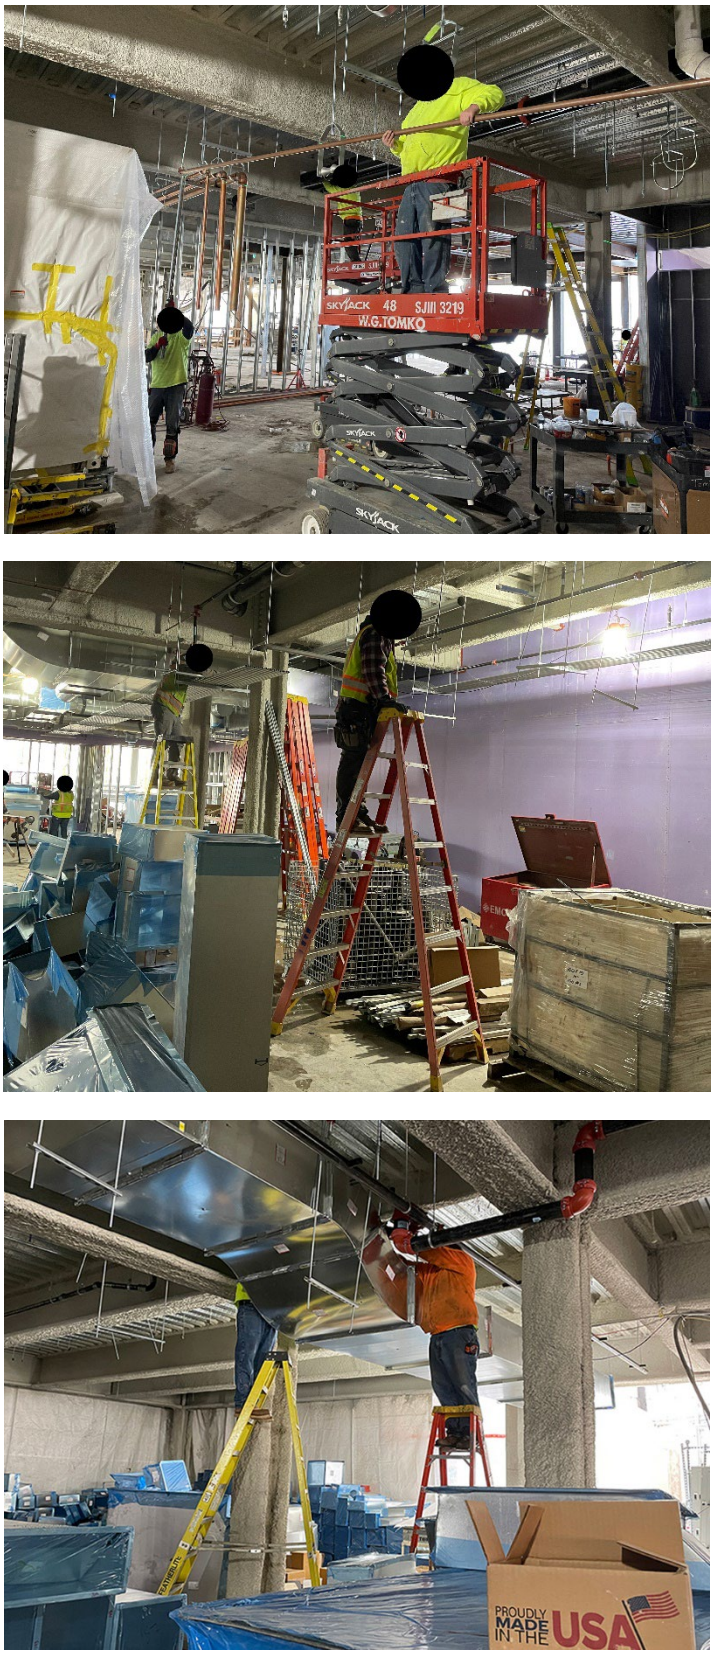


Figure 8. MEP workers working on lifts in more open zones and on ladders in more congested zones

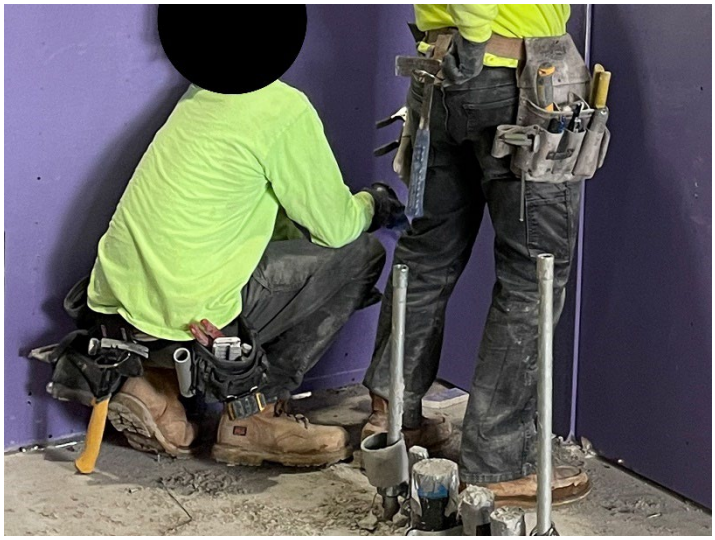


Figure 9. The tools carried by the drywall workers doing bottoms



Figure 10. Ironworker 2 demonstrating the wire tying move

manipulation problem (Figure 10). Recognizing this dexterous character of tying in my description, they shifted their attention to more physical activities in the task. Aside from my suggestion of robot tool delivery to Researcher 2, Researcher 1 also suggested that “bringing things from point A to point B” could be more useful in this task, which we later saw not viable as I previously noted.

Understanding tasks in this granular way and based on their qualitative character generated two specific inclinations in ReAC’s technology design work, displaying the group’s alternative conception of automation. First, we began to draw out the time- and skill-invested parts of construction work. Considering the project’s general vision, we sought systems that would keep these in the hands of tradespeople or assist the way they perform them. This also aligned with how some workers envisioned the role of robots as assistants. Second, this way of seeing work shifted our attention to tasks involving menial physical activities for support, which I call shortly the ancillary activities. These activities, such as garbage delivery in cleaning-housekeeping, concurrent cleaning in fireproofing, or

tool-hardware delivery in framing and carpentry formwork, involve a high degree of worker mobility to access points or transfer items. Automating these activities appeared more salient to us in terms of meaningful robot support. Three use cases that we conceived represent this conception. While I present tool-hardware delivery (to support carpentry formwork) that we developed into a scenario and a prototype, the other two of them, garbage delivery (to support cleaning-housekeeping) and concurrent cleaning (to support fireproofing) can be found in Appendix D.

Tool and hardware delivery activity appeared as a potential form of support that a robot can provide during my observations of framing and carpentry formwork. Framer 1 did rounds of walking between the rolling scaffold, where he kept backup hardware and cut tracks, and various installation points. During both track and stud installation, as he and Framer 2 advanced along a track together, he worked kneeling or bent down with a tool harness on him (Figure 11). As the supporting worker in formwork, Carpenter 2 was even more mobile between the workbench and installation points, passing Carpenter 1 clamps, bracings, bolts, and sometimes power tools as the worker advanced along a rebar cage and installed panels (Figure 11). Seeing both these delivery actions, I considered that the same concept system I presented for rebar tying and MEP ceiling installation, which would follow the workers and autonomously navigate between points, might be a useful instrument. I imagined that the robot would operate close to Framer 1 and carry the items that he keeps in his harness, such as a tape measure, snips, impact gun, or backup screws. In carpentry formwork, the robot would be loaded by Carpenter 2 and keep an adequate load of power tools, clamps, and bolts in proximity to Carpenter 1. By taking over a menial supporting activity while the workers focus on the skill-invested, dexterous parts of the work, such as track cutting, panel preparation, and installation, the system would fittingly represent robot accompaniment.

This form of support seemed credibly useful to us in these tasks than in rebar tying and MEP ceiling installations for a couple of reasons. One was the amount of tool and hardware delivery and items being passed between workers, as opposed to rebar tying, where workers use fewer materials and items, and MEP workflows, where workers keep their stuff with them on lifts. The system would presumably reduce the amount of back-and-forth and eliminate the tool harness. For the most part, workers were receptive to the idea as well. When I described the system and its potential benefit to Framer 1, he approvingly stated: “it sounds like a good idea, it really does.” Some skepticism followed, however: “but our job is so unpredictable… I don’t see how it [a robot] would know when you’re just walking away to go to the bathroom, walking away to go get a drink.” Carpenter 2’s



Figure 11. Framer 2 working alongside a track (left) and Carpenter 2 preparing to feed hardware (right)

feedback was simple but more positive: “if I’m walking less, it’s good,” he said during a chat on the site. The second reason why the idea seemed more credible was that, unlike MEP ceiling installation, these tasks take place in relatively less congested zones more suitable for effective robot navigation, with presumably fewer “freezing robot” encounters.[10] Further, an RL-based navigation-only system would be sufficient to facilitate this meaningful support. To determine whether these intuitions made sense, I picked carpentry formwork, considered the stronger case, and performed a deeper analysis.

This analysis and documentation helped us make a case for the use case and shape it further into a specific scenario, which, alongside the final system, displays the ReAC group’s reimagination of robots role in construction. That alternative role is handling ancillary activities as instruments of task support in construction. As a virtual testbed, the scenario tied the conceived tool and hardware delivery function and the technical implementation to the existing formwork workflow. It also helped to identify the robot navigation capabilities needed to perform the delivery function in a conversation with Researcher 2. She then evaluated and translated these into a detailed implementation scope.

The analysis and documentation conducted by Researcher 5 and I focused on Carpenters 1 and 2’s work zone, routines, and tools. Performed by these workers with lead and supporting roles, this task aimed to build formwork for a cast-in-place concrete wall stretching north-south on the eastern edge, as I described in the previous chapter. “Did you take a closer look at the tools?” asked Researcher 1 when we first discussed the tool and hardware delivery use case for this task. Similarly, Researcher 2 inquired: “how heavy are those [tools]? Can you just carry a bolt impact gun like that, safe?” These were important questions, urging us to understand the specificities and reassess whether the use case was as credible as we thought. Based on the time-lapse recording and photos, I delineated the approximately one-hundred-and-fifteen-foot long and twenty-foot deep, elongated rectangular zone, including the stations where the workers kept their tools, hardware, and panels as well as various obstacles (Figure 12). These include the workbench on the northern end carrying the power tools and small hardware like bolts, next to a large rectangular container full of clamps, and the main panel station in the middle holding panel stacks on pallets. These items are also among the obstacles for a robot. The other obstacles range from columns through leftover material piles and other trades’ material stacks to standing construction equipment such as ladders and tower lights. Photos of this scene and items also enabled me to document the shapes and dimensions of tools as well as the spacing between obstacles for navigation. This confirmed my initial sense that the work zone, unlike other extremely congested ones, was suitable for effective robot navigation.

In parallel, I conducted a detailed analysis of the workers’ routines through video coding, with a particular focus on Carpenter 2. Video coding went a step further than the on-site observations, capturing specific movement frequencies and trajectories. As I outlined in the previous chapter, Carpenter 1, as the lead worker, handled the formwork panel installation, while Carpenter 2 supported the installation by setting temporary panel stations close to the areas of installation along the cage, transferring the panels, and delivering tools and assembly hardware to Carpenter 1 (Figure 13). Carpenter 2 multi-tasked between these activities, which all involved varying intervals and walking distances. For instance, for every five to seven panels that Carpenter 1 installed, he shifted the temporary panel station, which was mobile on a pallet jack, or walked to the main station, took a fresh stack, and brought it to a new area of installation. Depending on how far the area was from the main station, the distance varied between five and fifty feet. For thirty seconds to a minute, he

10 “Freezing robot” problem is a robot not being able to move forward as a result of determining all possible paths in a congested setting lead to collision, based on Peter Trautman and Andreas Krause’s (2010) article. See more in Trautman and Krause (2010).

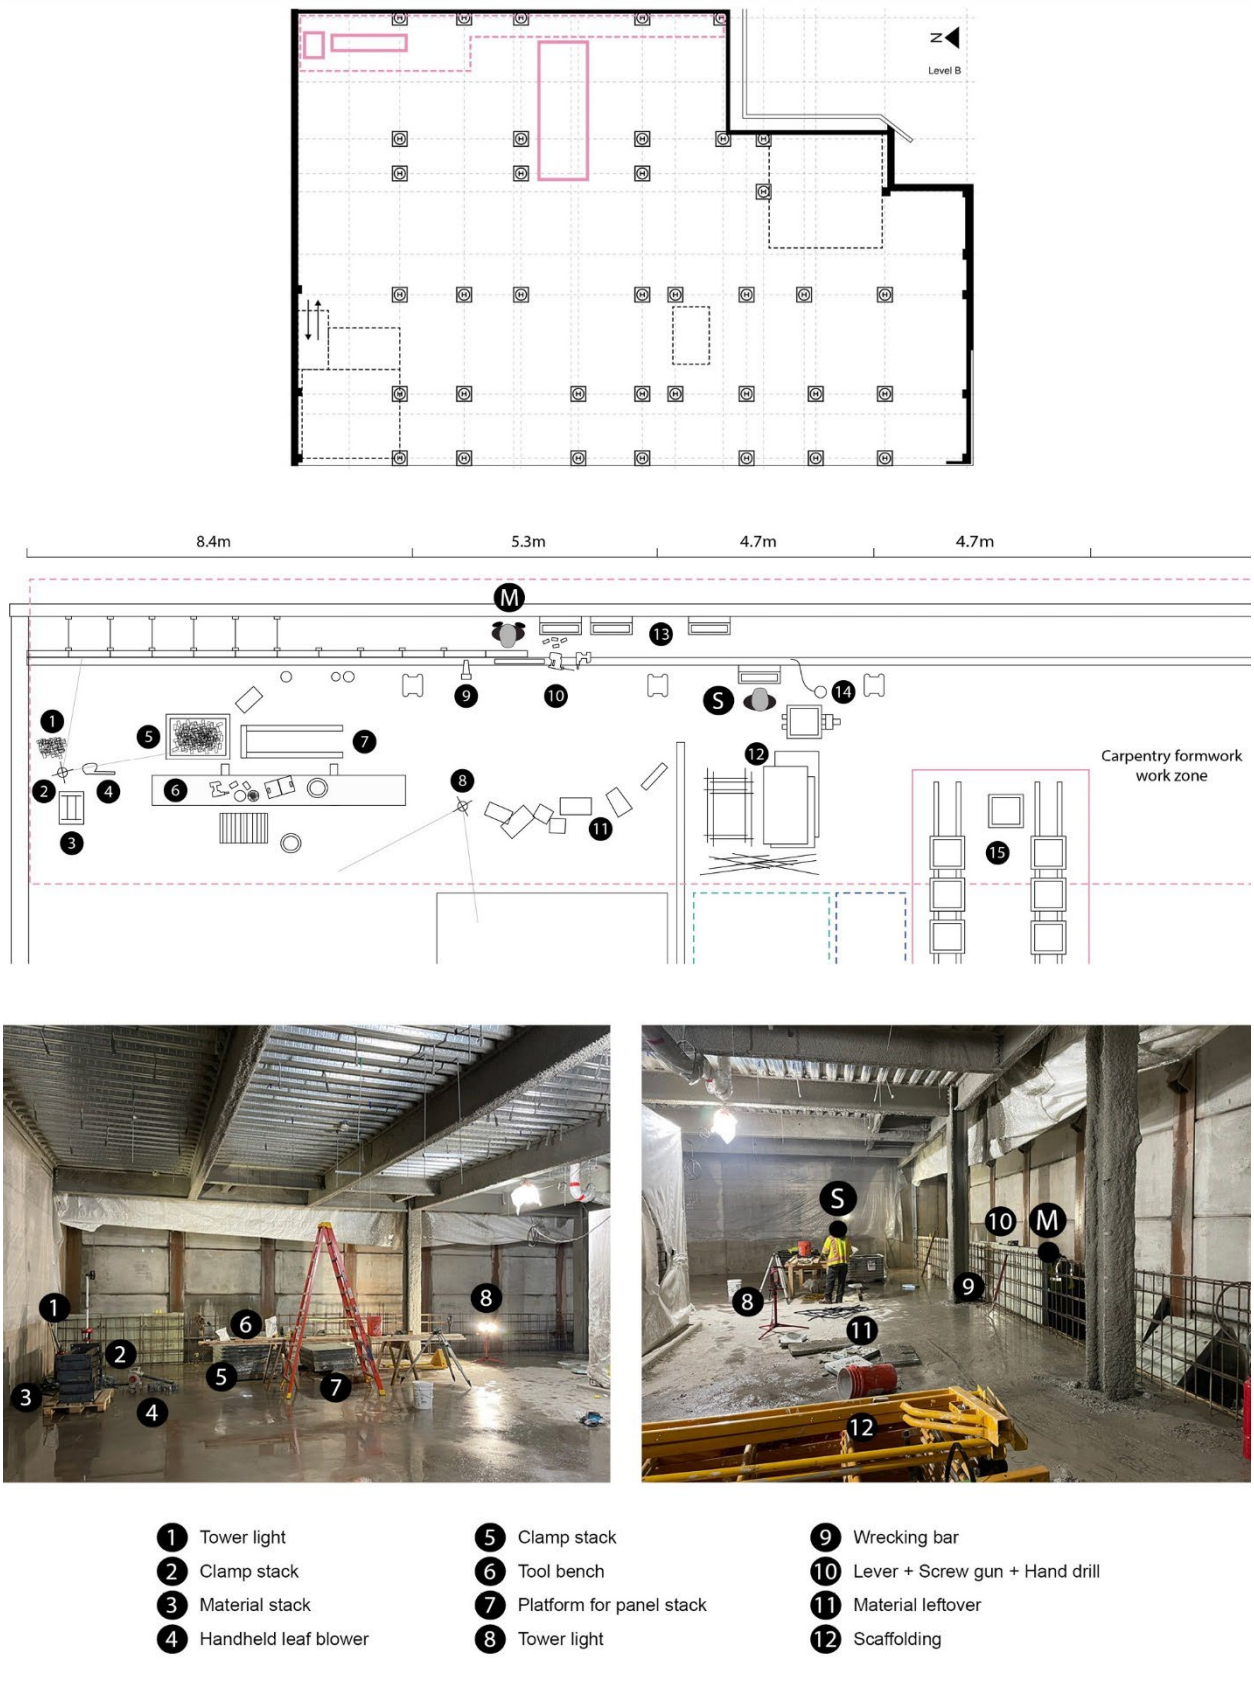


Figure 12. Carpentry formwork zone on Level B—main worker (M) and supporting worker (S)

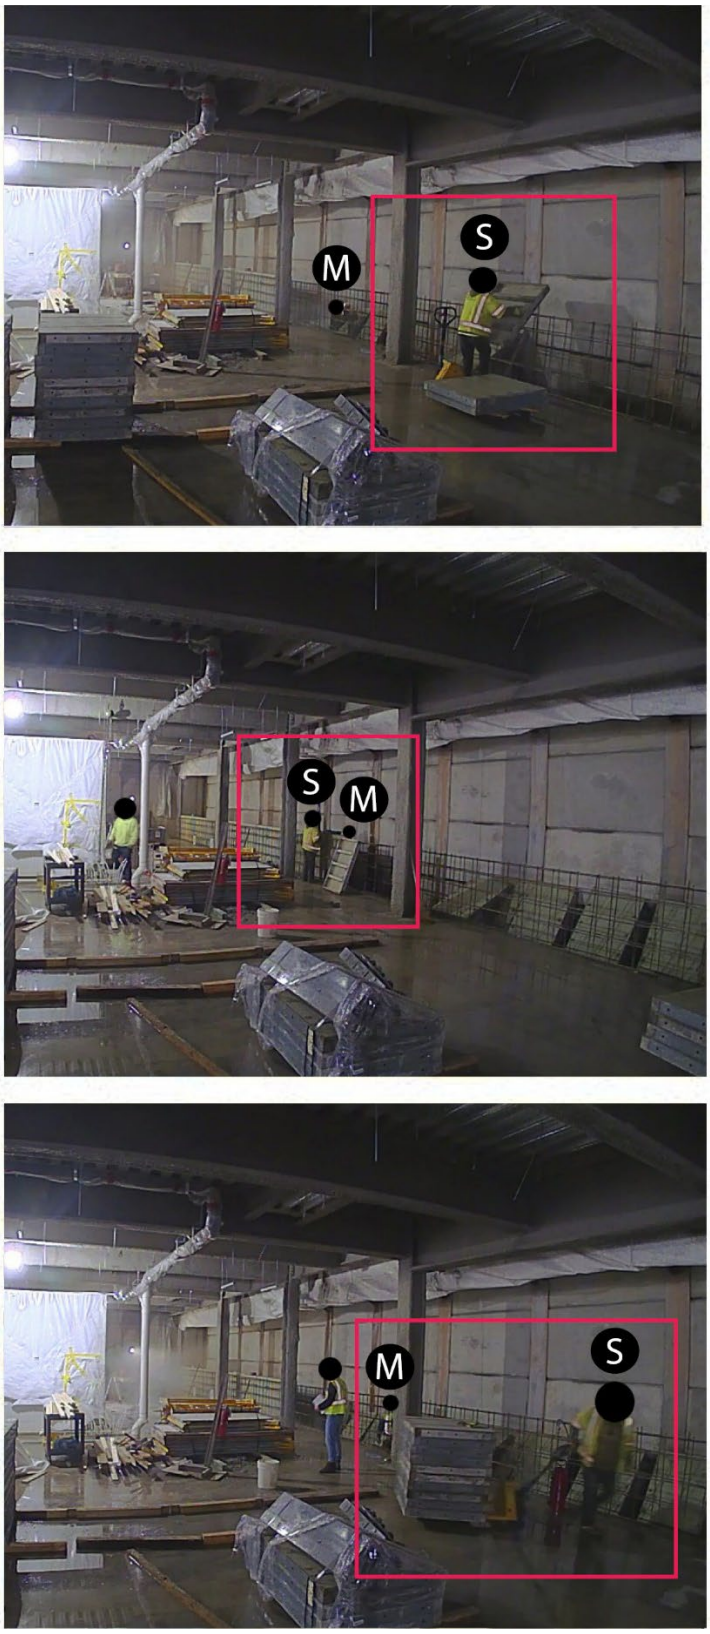


Figure 13. Carpenter 2 transferring the panels (top), delivering tools and hardware to Carpenter 1 (middle), and carrying a panel stack to an area of installation (bottom)

picked, transferred, and sprayed a panel. To do so, he took short but rapid walks between a temporary station and the cage. He delivered tools and hardware at irregular intervals. To deliver hardware, he walked from wherever he was to the workbench and from there to Carpenter 1 with clamps, bracings, or bolts in his hands. Depending on how far Carpenter 1 worked from the workbench, the distance could reach more than fifty feet. To deliver tools, these walks occurred between where a particular tool was left and where Carpenter 1 had already advanced, which could be varying distances.

A series of key observations in this repetitive course of action allowed me to make a case for the conceived tool and hardware delivery system. First, the three supporting activities, which might seem trivial individually, amount to a heavy physical load when performed for hours in the intervals I mentioned. Second, tool and hardware delivery, which is probably the dullest and the least important of the three, is repeated more than the other two. While Carpenter 2 performed panel station setup two to three times and the panel transfer thirteen times an hour, he performed the tool and hardware delivery nineteen times within the same period. Because the activity was contingent on Carpenter 1's needs, it reoccurred erratically. Since the distance depended on Carpenter 1's location, each delivery required a longer walk as Carpenter 2 advanced along the cage, increasingly farther from the workbench or other hardware locations. These facts substantiated my initial belief that by keeping a substantial tool and hardware load in proximity to the lead worker and reducing the number of deliveries by the supporting worker, the conceived RL-based navigation-only system could be meaningfully supportive. The third key observation concerned the number of tools delivered, which confirmed that the concept system could be of greater value than in rebar tying and MEP ceiling installation. Carpenter 2 fed clamps, bolts, and bracings to Carpenter 1, in addition to occasionally bringing his power tools, including a drill, impact gun, level, and angle grinder, from where they were used last to his current location (Figure 14).

To virtually test how the conceived system might work within the existing workflow, I developed a detailed scenario. According to this scenario, the workers can prompt the robot with a destination point from a pre-recorded selection, and the robot, functioning as a mobile supply station, autonomously navigates to the destination by avoiding collision with workers and obstacles. For instance, in a workflow where the lead worker is installing and the supporting worker is ten-to-fifteen steps ahead queuing up new panels, when the lead worker runs out of hardware, any of the workers may send the robot to the workbench. While the robot figures out the way, the supporting worker may continue the panel activities without pausing. Once the robot arrives at the workbench, the supporting worker goes to it, loads it, and sends it back to the lead worker. While the supporting worker quickly returns to his position, the robot can find its way back to where the lead worker stands. The robot holds more supply than the supporting worker could deliver in one round, so it will take more panels and time for the lead worker to run out again. This not only eliminates rounds but also enables separating the supporting worker from the lead worker for a longer period, allowing the supporting worker to maintain their pace in panel station setup and transfer. After installing three panels, the lead worker shifts position and can prompt the robot with the supply to move to that new position. When the workers migrate within the zone with other tools and a pallet jack of panels, the robot will again be helpful. In a variation on this scenario, the workers can define themselves as a goal position for the robot to follow in real time and keep the tools and hardware in proximity while maintaining distance and avoiding collisions.

Thinking through this scenario, I proposed two types of robot navigation needed to perform the delivery function to Researcher 2. The first one is the capability to safely navigate in an indoor, unstructured, and dynamically changing carpentry work zone from any starting position A to any goal

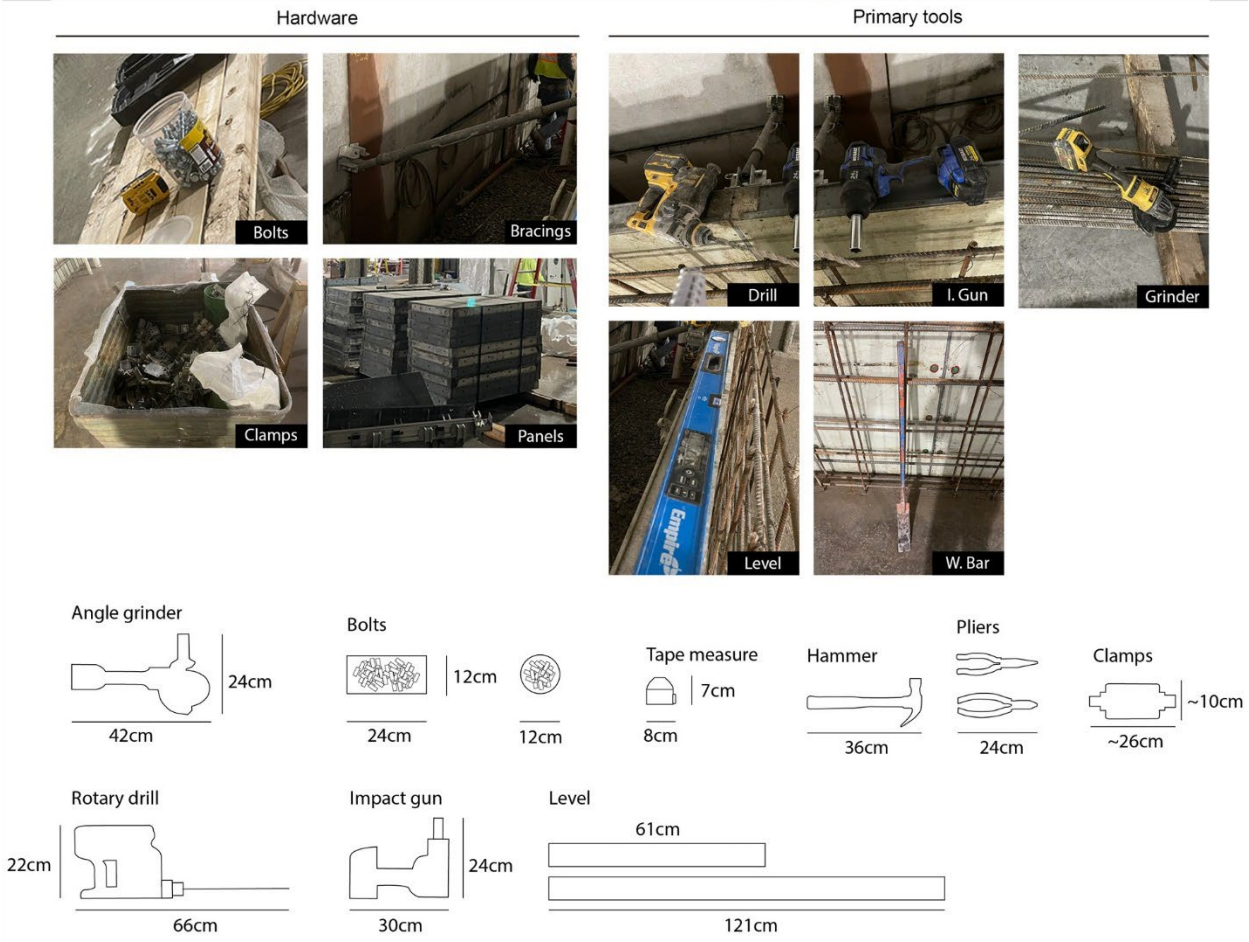


Figure 14. Tools and hardware used during formwork panel installation

position B without colliding with workers or obstacles. This would be sufficient to successfully perform the function. The second one is the additional capability to follow a worker as the goal position without breaking safe distance or colliding with workers or obstacles. These capabilities sound simple in description, but, as Researcher 2 and many other robotics researchers know, they are challenging to develop with a software-hardware integrated system. Researcher 2 identified the first capability as the baseline for achieving the function and a realistic goal to achieve for the project.

From Site to System Development: Practical Inputs

In addition to breeding imagination and conceptual guides in ideation and shaping partly the way we view work and situate a robot, the New Scaife Hall site study provided a set of practical inputs to the development of the final tool and hardware delivery system for formwork.[11] These inputs related to the knowledge of specific site conditions, obstacles, and density of people helped to select an appropriate chassis for the system, fine-tune the sensing of the system, inform the simulations used to test the system, and design an on-site study to evaluate the system. Through brief accounts of

11 For details of the system and its development, see Yuning Wu (2024).

these efforts in this section, I foreground this third role, informing the technical development, that the ethnographic study played in system development.

Chassis selection: Knowledge of specific site conditions in addition to technical considerations governed the chassis selection. Technical considerations such as reliability and hardware compatibility were the priority for an initial period. Led by Researcher 2, Researchers 3, 4, and 6 tried and failed with different mobile bases. For about five months, they worked with a used Fetch Mobile Manipulator until it suffered an internal computer failure.[12] Although a second-hand Clearpath Jackal was considered a potential substitute, a TurtleBot temporarily helped carry the ongoing development efforts, including a preliminary demo for the main project sponsor.[13] Instead of the substitute under consideration, a custom-built base named Rocky (built at CMU) was brought to the lab, but this robot, which we all thought would be the final fix, was soon dropped due to its risky quick acceleration. Practicality was the quality that I sought in these options. More precisely, the chassis should be able to withstand the site conditions, and none of these bases seemed to check that box. Even an extension cable would possibly challenge Fetch and Rocky, which were tall, low, and with small wheels. Jackal, which we never acquired, was the closest prospect but still slightly too small.

Our attention shifted somewhat from technical considerations to site conditions when we received Rocky. “What will the robot look like?,” was the question Researcher 1 asked about Rocky, bringing chassis design into our conversation for the first time. Because the chassis was low and bare without any casing or sensors placed, he wanted to think of design adjustments to modify it for the site. Ultimately, our chassis during formwork should be able to roll reliably on a concrete slab, which is relatively flat but in a messier condition than outdoors. It should have a robust body and thick rubber tires for rolling over dust, puddles, cables, hoses, and small garbage, which are never missing on the floor, including in formwork zones. Rocky was still not a perfect fit but what we had to work with. It would require many adjustments. Luckily, after Rocky’s breakdown, with the knowledge of these conditions in mind, we acquired a second-hand Clearpath Husky, a true field robot much better suited to actual site conditions.[14] To prepare the robot for experiments, Researchers 2 and 3 equipped it with a standard sensor suite for navigation, including a 2D LiDAR, a 3D LiDAR, an IMU, and two RGB cameras to capture sensing data, along with a laptop with sufficient CPU-GPU power for data processing (Figure 15). Later in the project, they designed and built a mount structure and a sensor box, keeping the 3D LiDAR and IMU together at an appropriate height (Figure 15).

While the knowledge of physical conditions informed the chassis selection, the documentation of formwork tools and hardware guided the modification of the final chassis as a tool and hardware carrier. I kicked off this effort with the conceptual design of a container, taking into account a couple of constraints. For instance, the onboard computer would sit on the back half of the chassis, so the container had to sit on the front half. An ideal maximum size could be twenty-six inches long, sixteen inches deep, and twelve inches high, which would safely fit without protruding or blocking the 3D LiDAR and allow enough space for the carpenters’ bare minimum kit, including the impact wrench, rotary drill, angle grinder, clamps, and bolts (Figure 15). The amount the container could hold would not exceed Husky’s payload. This design had to be modified significantly to work well with the sensor placements. For this, Researchers 2 and 3 adjusted the size, reworked the positioning, and, additionally, executed an assembly that made the container retractable (Figure 15).

12 See “Fetch Mobile Manipulator” by Melonee Wise and colleagues (2016) at Fetch Robotics.
13 See “Jackal” launched in 2014 on Clearpath Robotics’ (2023) website, and “TurtleBot” on TurtleBot’s (2023) website.
14 See “Husky” launched in 2011 on Clearpath Robotics (2023) website.

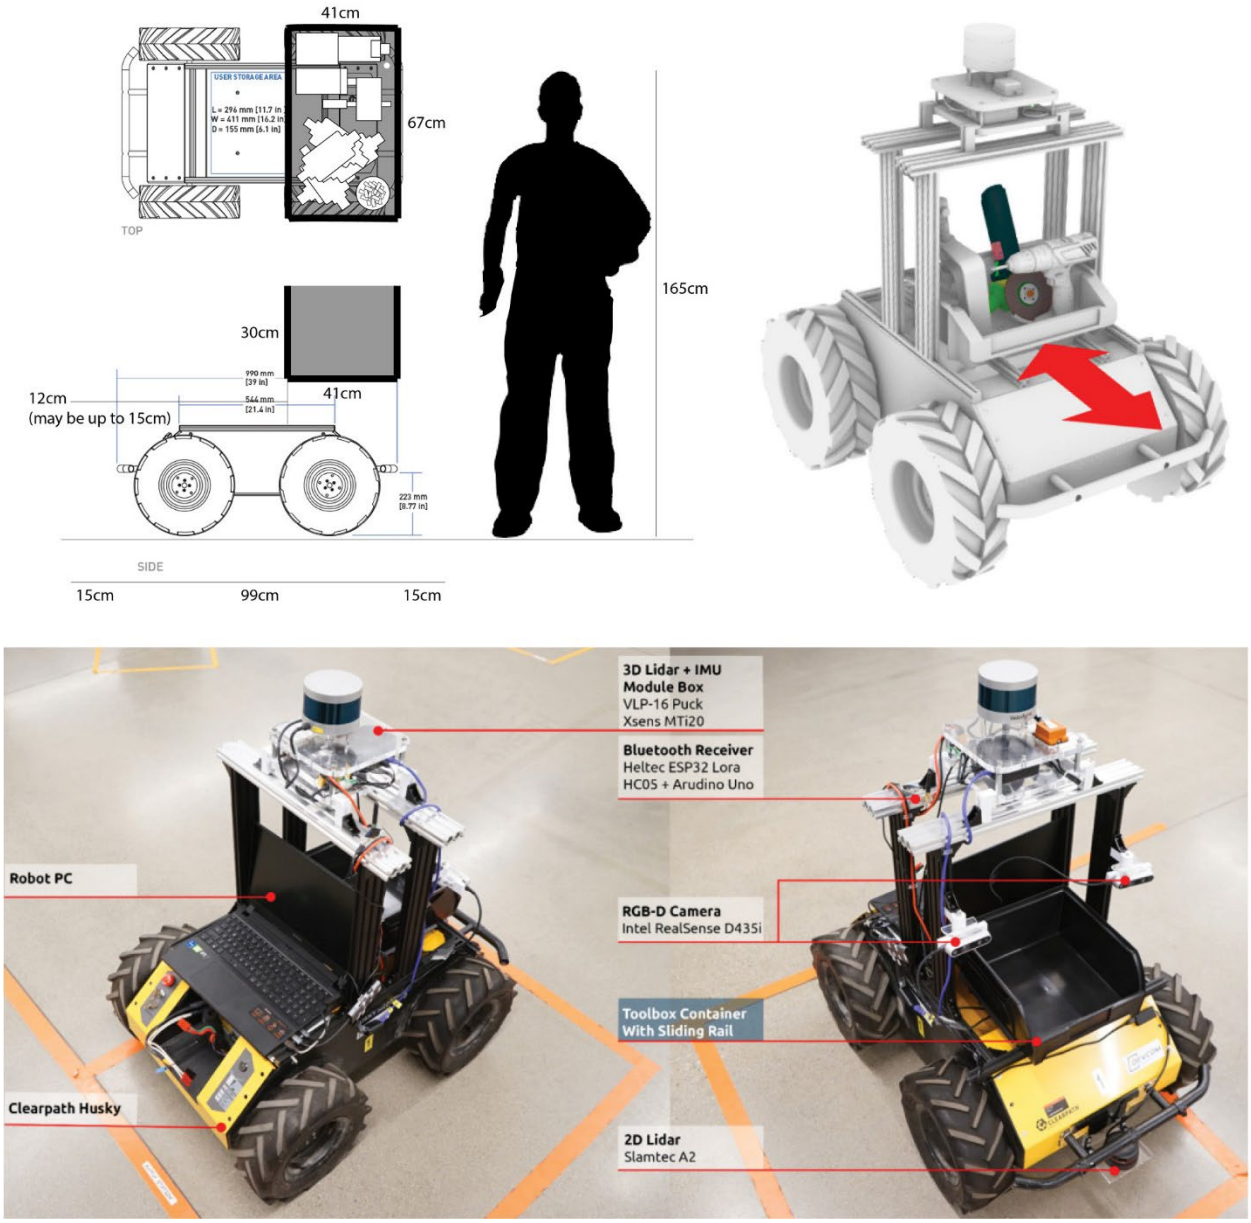


Figure 15. Initial concept design for container (top left), Researcher 3's final design (top right), and robot in final shape with the sensors and container (bottom) (ReAC Research Group / Wu & Wei, 2023)

Fine-tuning sensing: Specific obstacles I documented on the site were an informative part of the technical work on sensing. Work zones in Level B had static obstacles that were impossible or difficult to move, such as columns, frames, pipe heads, and material stacks, as well as dynamic obstacles that workers moved around, such as ladders, workbenches, gang boxes, and garbage piles. Seeing site scenes along with a catalog of these objects that I prepared with Researcher 5 gave Researcher 2 and others a much more vivid sense of the specificity of this particular setting compared to those she considered "well-studied," such as warehouses or outdoor traffic. In a later presentation, she identified a construction scene as "unstructured," with more "complicated" obstacles of diverse and complex shapes, and, thus, "unknown territory" in robotics. Safe navigation

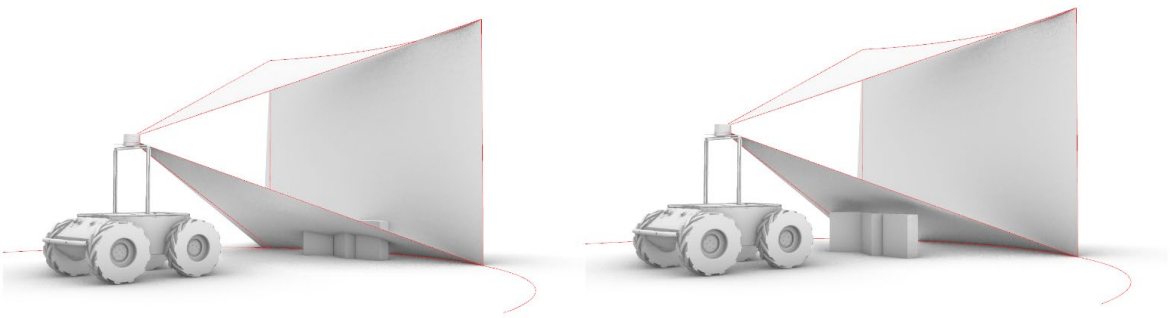
in a work zone in one such setting required our system to detect both the static and dynamic obstacles Researcher 5 and I exemplified.

To achieve this, Researcher 2 took what I see as a layered approach. “We want to map [the] static obstacles,” she said, beginning to unpack the first layer for us, in an early technical update session. To capture the obstacles at different heights, she, together with Researcher 3, combined 2D mapping to generate a 2D occupancy grid map of a floor and 3D mapping to reconstruct its 3D point cloud and estimate the robot’s position.[15] Based on the current position, the map including the static obstacles, and a goal position, the global path planner generates the trajectory for the robot. “[However,] we don’t actually rely on global path planning to be the final solution,” she continued, and explained the second layer: local path planning would be where the system “recognizes what is really around,” including items that are dynamic or not captured in the map, through real-time LiDAR data, and “correct[s] itself.” With the granular information that she had about the objects on the site, Researcher 2 also repeatedly pointed out “low obstacles,” such as pipes, at a short distance as well as “corner cases,” such as assorted garbage piles, that the LiDAR could possibly miss. A sense of low object heights on the site helped Researchers 2 and 3 fix the 2D and 3D LiDAR placement to minimize those misses (Figure 16).

Other than detecting objects, safe navigation in a work zone also required detecting workers and tracking their walking trajectories when working alongside the robot. The input that helped Researcher 2 and others in implementation was the worker–robot interactions that the on-site observations and the scenario helped anticipate. More precisely, these were potential close interactions with carpenters working directly with the robot and less likely but possible close encounters with other workers crossing the robot’s path. While searching for a model, Researcher 2’s concerns were computational efficiency because the system would rely on a single onboard computer and accuracy. To find the best fit, she supervised Researchers 7, 8, 9, and 10 in running quick tests to evaluate a list of camera- and LiDAR-based models for both detection and tracking. A strong candidate was “YOLO,” a camera-based detection model faster at processing frames but whose accuracy depends on the clarity of faces, as Researcher 7 informed the group while comparing it to other models.[16] This would be an issue when workers’ heads may exit the frame of the robot’s camera while they pick up hardware or pass by, which Researcher 2 found out when she re-tested and implemented the model’s newer version. To maintain accuracy and, therefore, worker safety in these instances, she integrated an additional LiDAR-based model that detects legs, so that the two models, in her words, would “compensate for each other.”

Informing simulations: The documentation of obstacles and workers on Level B provided important insights for achieving a close representation of the formwork zone in simulations to test collision avoidance. A key technical concept that Researcher 2 brought to the project was “social robot navigation,” which refers to an area of research towards robot adaptiveness to dynamic human presence during navigation.[17] “Distinguish[ing] pedestrians from other obstacles,” as Researcher 2 put it, is important for social navigation in construction and currently a limitation of “one-step canonical” and “one-step RL” collision avoidance approaches. She followed a “two-step RL” approach, in which she combined a one-step canonical approach to prevent collisions with static obstacles and RL to prevent collisions with workers. In this way, as the global path planner “defines a trajectory” and “divides it into segments,” for each segment, the local path planner, which

15 3D mapping was done through “LOAM (LiDAR odometry and mapping).” See more on LOAM in Ji Zhang and Sanjiv Singh (2014).
16 See YOLO in Glenn Jocher and collaborators (2020).
17 For more on “social robot navigation,” see Christoforos Mavrogiannis and coauthors (2023).



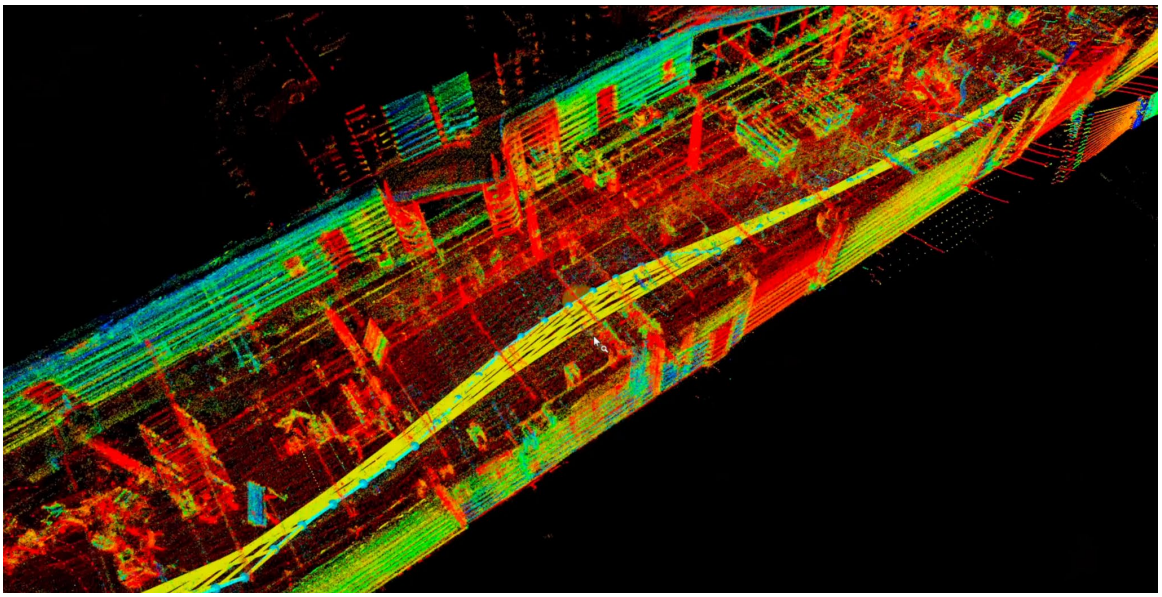


Figure 16. Researchers 2 and 3's 2D-3D LiDAR combination to minimize missing low obstacles at a short distance (top) and 3D point cloud reconstruction of lab space as part of a test (ReAC Research Group / Wu & Wei, 2022)

combines the two steps, first generates a set of obstacle-free "directions" and then finds the "most socially compliant" one, as Researcher 2 explained in her first introduction of this approach in a meeting (Figure 17).

In the same meeting, Researcher 2 showed some initial test results on a CMU-developed simulator, called the "Social Navigation Simulator," which she and Researcher 6 were using to evaluate the selected RL algorithms.[18] The simulation scenes I saw depicted an agent walking through a tight crowd of others. In contrast, a work zone would usually feature two to three workers, who would maintain distance from others to avoid conflicts and safety risks, as I described in the previous chapter. Overlaps between zones might occasionally create instances similar to those in the simulation, in which workers would still limit proximity. "It [the density of workers] is quite different from… the density of pedestrians in public space," I opened the conversation, adding, "it is a little bit more sparse." Researcher 2 was on the same page: "so maybe we should test fewer agents, but

18 See "Social Navigation Simulator (SNS)" by CMU RoBot Intelligence Group (BIG).

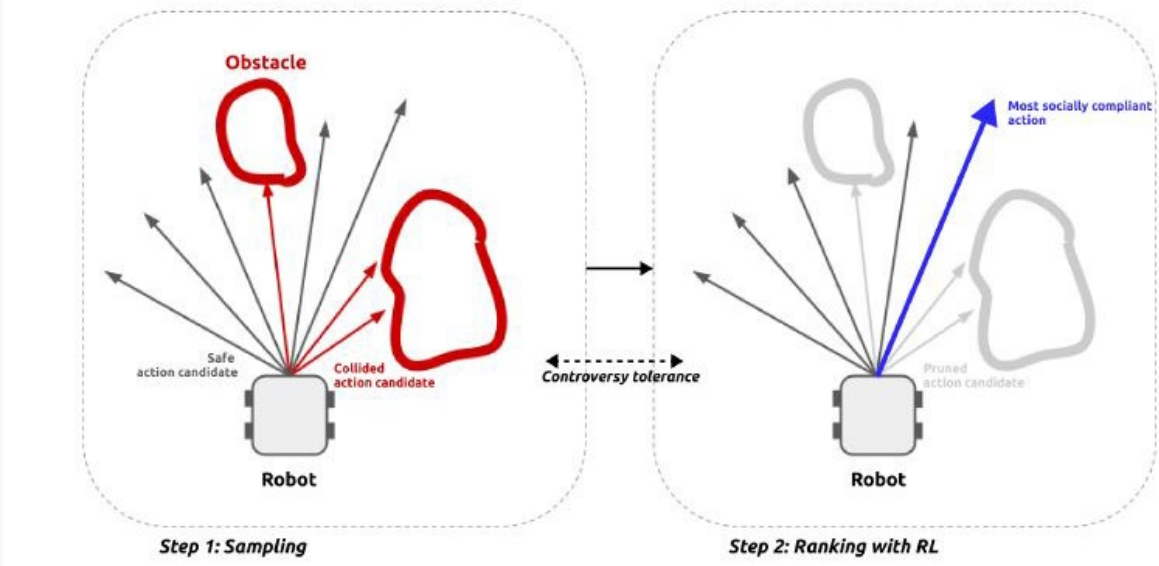


Figure 17. Researcher 2's two-step RL approach (ReAC Research Group / Wu, 2022)

more edgy collision instead of marching into ten people." Ultimately, she trained the final model with a scenario that she called "more towards ours," including four to eight agents. She also tested it for possible close encounters by recreating those in the simulation, a worker walking across, a worker performing an activity nearby, and five workers moving around the robot (Figure 18). The site and scenario also informed the physical sandbox setup that Researchers 2 and 3 put together. The documented shapes and dimensions of objects and the dimensions of movement paths in the formwork zone helped them build a similar configuration consisting of high and low obstacles or their imitations, in which the robot was tested with us initiating close encounters (Figure 18).

Designing on-site evaluation: Finally, the scenario for situating the system also helped in designing an on-site evaluation with construction workers. The aim was to evaluate the performance and usability of the robot as well as the usability of its UI, a keypad controller that Researcher 3 designed and implemented (Figure 19). She focused on creating a simple design with buttons for quick and easy use by carpenters who multitask with their hands during formwork, at times wearing gloves. Through the interface, carpenters can record a set of points where they want the robot to go to load or switch the tools and hardware throughout their workflow. For instance, they may set a point every three panels, and the robot can shift with the lead worker in those intervals. They can then send or call the robot back between these points. For example, the supporting worker can send the robot to the loading point and, after loading it, to a point where the lead worker is installing. For safety, workers can check the navigation status and cancel the robot's action in real time. For the on-site evaluation, Researchers 2 and 3 and I designated a test area approximating a formwork zone, and we asked worker pairs to enact a workflow approximating the one I observed on Level B, with our robot as part of the workflow. Specifically, we asked them to perform three sets of interactions we called prompts, two of which represented the key moments in our scenario.

The first prompt from the scenario involved using the robot between a tool and hardware station, representing the workbench in the scenario, and an area of installation (Figure 19). Similar to the scenario, we observed a worker assigned the supporting role calling the loaded robot from the station via the controller to the area. The robot successfully navigated to a pre-recorded point close to its co-worker in the lead role without any collisions. The robot remained at that point with a load of

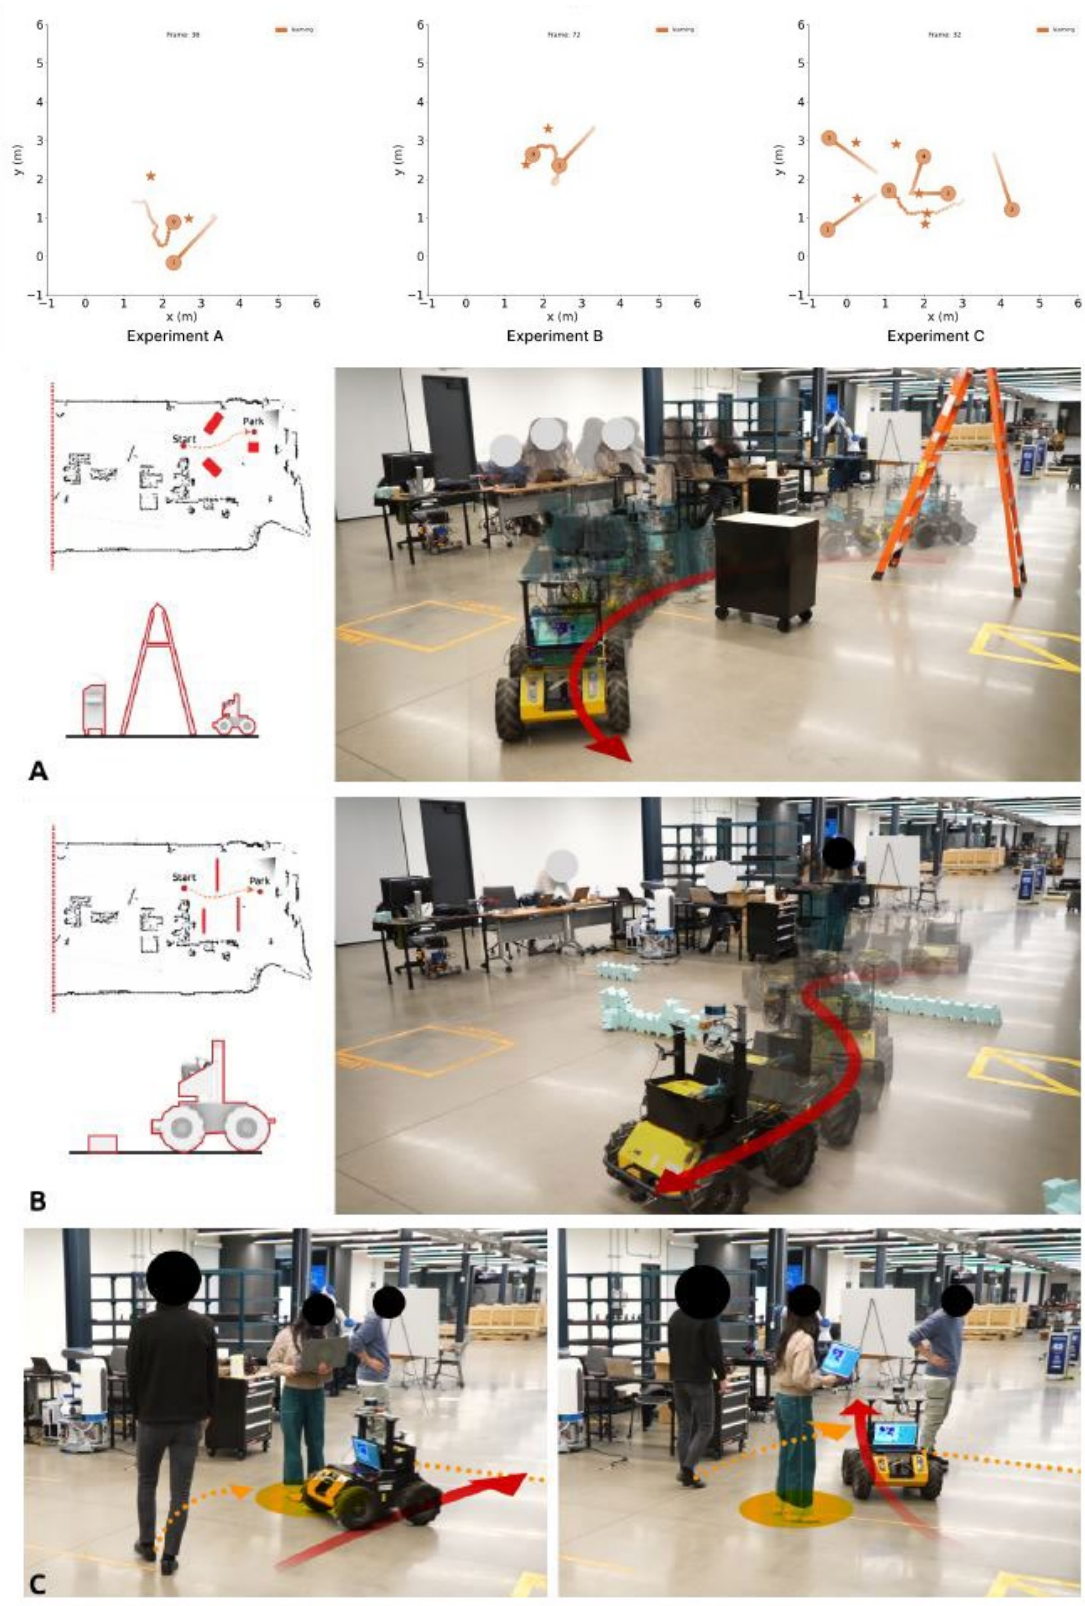


Figure 18. Researcher 2's tests for possible encounters in simulation (top), and Researcher 2 and 3's tests for high and low obstacles and close encounters in a physical sandbox setup (middle and bottom) (ReAC Research Group / Wu & Wei, 2023)

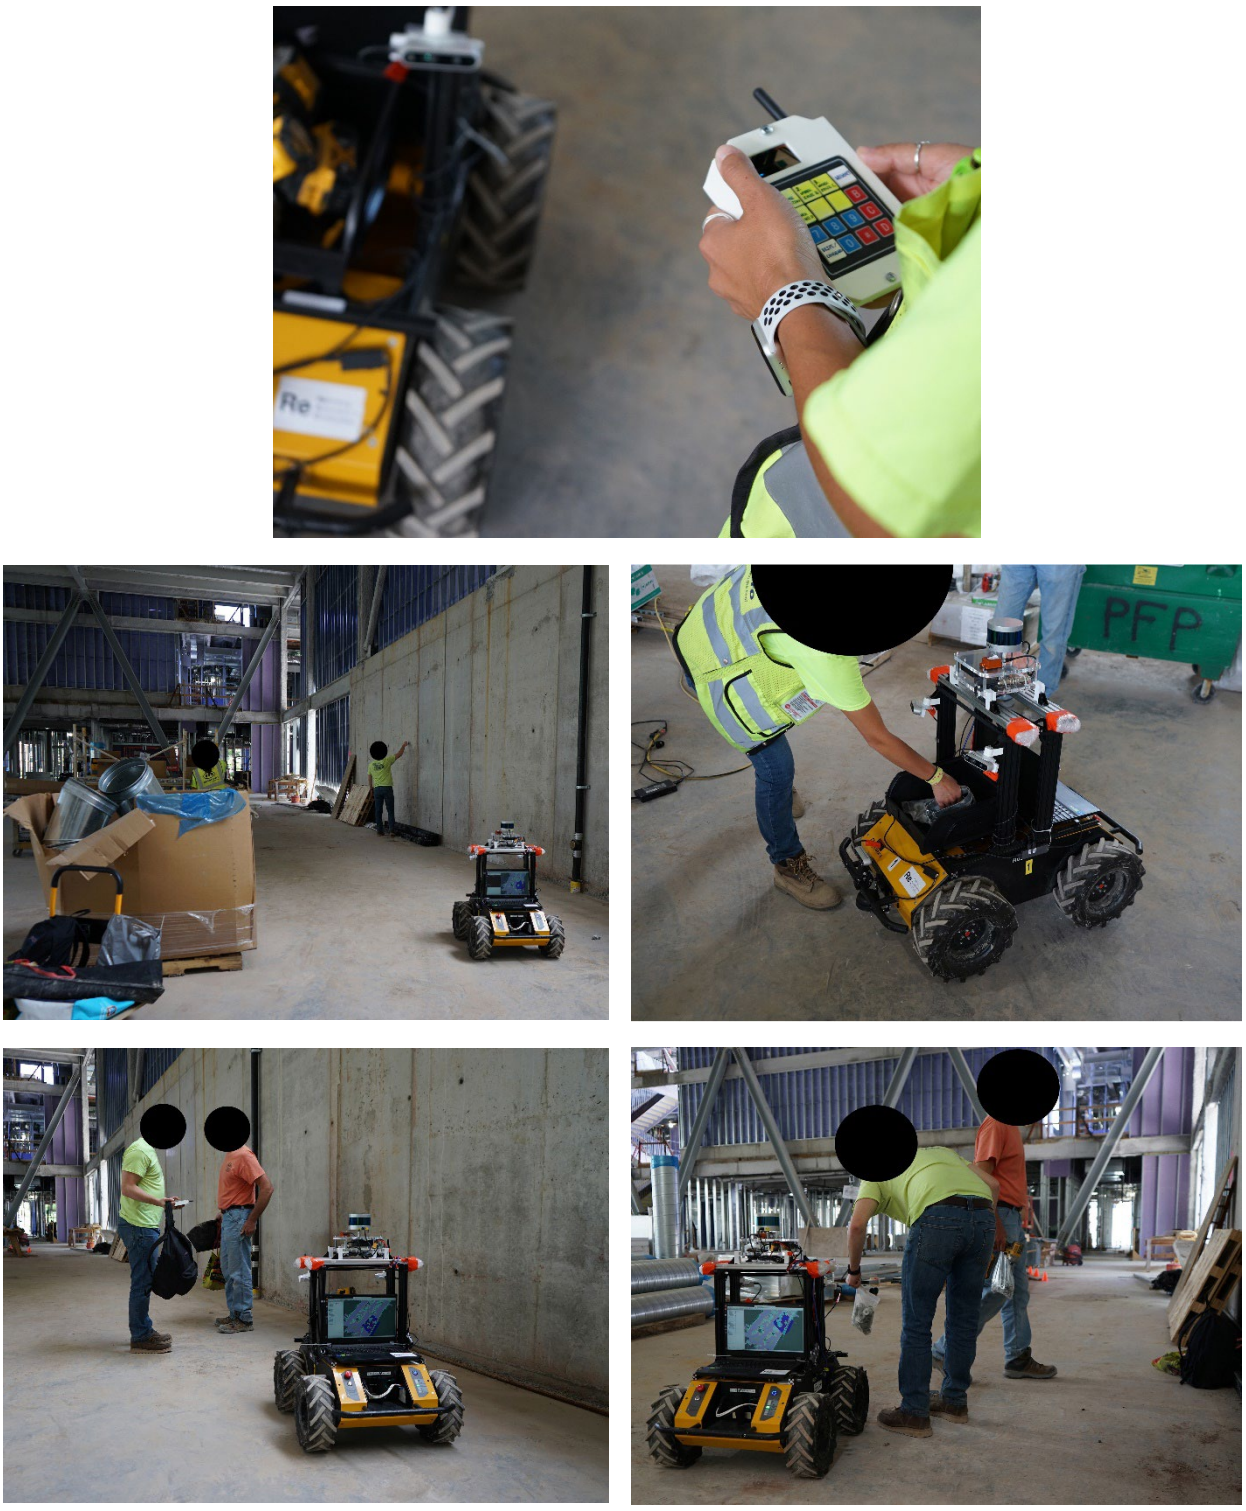


Figure 19. The keypad controller for the robot (top), participants (1) using the robot between the tool and hardware station and an area of installation (middle row), and (2) while shifting points (bottom row)

hardware at the lead worker's disposal for a while. Then, when we assumed that the load had been fully used, we asked the supporting worker to send the robot to the station. The robot took a similar route back while the workers continued their activities. Once the robot arrived, the supporting worker

attended to it briefly at the station, reloaded it, and sent it back. The worker could immediately return to their activities as the robot made its way to the lead worker. Other than the required brief attendance to load the robot, the supporting worker did not have to pause their other activities, which is what we expected based on the scenario. Although we did not have the opportunity to keep the participants continuing for an hour, even one short round implied that keeping a substantial load at their disposal would reduce the frequent delivery calls.

The second prompt from the scenario involved using the robot while moving between points in the zone (Figure 19). The workers started by imitating their respective activities. Then, when we considered that they should be done and ready to transition, we watched the supporting worker send the robot to a pre-recorded point in the next area of installation. While the workers led the way by carrying other stuff, which could be large-size tools or panels on a pallet jack in the actual scenario, the robot with the tools and hardware was heading to the same destination. The robot successfully arrived at the point, and the installation work in this new area started. This shorter instance showed that the robot could also successfully work as a carrier, sharing the load with the workers during transitions, as imagined in the scenario.

4.3 Summary

By incorporating an ethnographic study of construction work into system development, the ReAC group implemented a rare and important interdisciplinary convergence of researchers with various skill sets, workers from various trades, and a construction site as part of a practice-based technology design effort. Rather than targeting a task and researching ways to automate, we went bottom up and sought ideas with multiple tasks (that we studied closely) in mind.

By looking into these efforts in this chapter, I showed how ethnography draws robotics and construction closer by identifying the ways the elements of the New Scaife Hall study guided our eventual conception and development of the tool-hardware delivery system. As our discussions that I presented in the chapter exemplify, the task descriptions and the data viewer helped immerse ourselves in the workflows, and collectively imagine systems in a more detailed and grounded manner. As we learned more about construction, our growing knowledge of tasks and existing knowledge of computing and robotics converge in ideating and evaluating systems. During those moments, our conceptual and technical views informed and balanced each other.

The division of labor in tasks that I conveyed in descriptions led to the connection made by Researcher 2 between accompaniment and the existing supporting roles on the site. This concept framed the way we sought similar roles for robots, and the way we ultimately demonstrated it with our final scenario and system. The knowledge of tasks' specific qualities that I broke down in the descriptions led to the emergence of our guardrails (specificity, credibility, and technical feasibility) with which we considered potentials and concept systems, as I traced in our conversations. Accompanying labor in a specific task in a credible way with a technically achievable system became our goal, based on which we discussed and discarded ideas.

Further, the knowledge of zones and interactions, the task descriptions, and the worker perspectives guided us while searching for and evaluating use cases as I discussed in the chapter. The challenges created by material zones for floor-wide robot navigation, the spatial organization of tasks in work zones, and limited interaction between these zones made us focus on supporting a typical two-to-three-worker unit focusing on one task in one specific zone. Workers' points on the situated nature of their tasks aligned with our aspiration to develop a system that fits into, rather than

disrupts, a task's natural course. Their reflections on the complexity of replicating skills with machines and concerns about deskilling made us focus on supportive than replicative system capabilities. These points as well as workers' direct feedback helped us evaluate and eliminate use cases. We conceived the attributes of a system that the workers view as suitable and useful through these points and feedback.

Importantly, the distinction of qualities in task descriptions helped us form a nuanced understanding of construction work, which lay beyond brief prepackaged descriptions served by or for technology designers. As I pointed out in the previous chapter, the descriptions conveyed that building tasks are difficult to encompass in a set of adjectives, such as "physical," "repetitive," or "mobile." The descriptions gave us the chance to comprehend how and to what degree each task embodies these qualities and, rather than what each task is, how it is actually enacted in practice. I showed that, based on this view of tasks, our analysis focused on the activities rather than tasks as we searched use cases or conceived potential systems. Rather than the general character of tasks, commonalities between tasks, or merely the basic repetitive activities in tasks, the group considered various activities in each task based on their qualitative character. As I discussed in the chapter, this shows our conception of construction as a richly layered practice, and we recognized the parts in which workers invest training time and skills and the intellectual parts to which they attach value.

This granular view of construction work made the group's focus on potential systems that would keep these in the hands of tradespeople or assist the way they are performed. Thus, we ideated based on a distinction between what might be suited to robots and what should remain with workers. Consequently, it diverted our attention from primary building activities, such as drilling, fastening, spraying, and installation, to ancillary ones such as garbage delivery, cleaning debris, rebar positioning, or tool-hardware delivery, which are only visible on site. These moves display the group's alternative approach to the historically prevalent task automation ambitions in construction, in which ancillary activities are objects of interest, and, as relevant alternative spaces to explore, deserve more attention. The idea is that fully or partially automating these activities may help reduce workers' load while preserving the time-invested, skill-invested, and valued parts of their work. We ended up exploring several of these, and finally conceived the final tool-hardware delivery system for carpentry formwork and situated the system within the task's existing workflow through a scenario.

The scenario and the system shows our reimagination of the role of robots, which are often conceived as instruments of task automation, as instruments of task support in construction. In addition to successful technical implementation, the system displays an alternative positioning of robots in construction, differing from that of single-task robots designed to perform primary activities. To determine how effective and credible this positioning might be, we again involved workers to evaluate the system. Ultimately, worker participants responded positively regarding ease of use and the way the robot performed its designated function. They questioned the fact that the robot "can't load itself," as one participant said, so the supporting worker has to handle it. When I explained that the robot takes away one activity off the supporting worker's plate and reduces the number of deliveries by keeping a substantial load near the lead worker, they grasped and appreciated the system's benefit more. Participants also offered further ideas, such as a chassis suitable for carrying materials, tooling to use the same chassis flexibly for multiple tasks, and a UI on a wearable device, and identified other carpentry tasks, including furniture assembly, door installation, and hardwood flooring, as suitable for the system.

Chapter 5

Construction Work Informing a System: Canvas

5.1 Outline

In the summer of 2022, I conducted an ethnographic study at the San Francisco-based construction robotics practice Canvas, in which I took a close look at its drywall finishing robot, Basil, and the ongoing development and deployment work around it. Through observations of everyday work activities (meetings, tests, and a construction job site deployment), interviews with the teams that develop, test, and operate the machine, an examination of its components, and the tracing of interactions with and around it via two prototype research tools I built, I explored the conceptions of drywall finishing work that come into play in the development of Basil, how these conceptions shape the development of the machine, and how the machine reconfigures drywall finishing work.

This chapter and the next broadly examines the early- and late-stage development and deployment work around Basil through the lens of these questions. In this chapter, I first introduce drywall finishing knowledge, which I traced the engineers at Canvas engaged during early design efforts that originated Basil and the typical workflow with the machine, and late-stage efforts to evaluate its performance. I outline the steps of the workflow, the tools that are used, the notion of "levels" in finishing, and the typical division of labor when the workflow is performed on a job site. By analyzing this knowledge of the task, which was both verbally transferred and demonstrated to me, I delineate the skills that the finishers develop and use in each step and discuss them according to two categories: explicit and tacit knowledge (Polanyi, 1966a).[1]

Then I trace and describe the two main lines of effort from which Basil originated in the past: one concerned the learning and examination of the steps of conventional finishing workflow through the lens of human and machine capabilities and the other focused on the operationalization of finishing quality. Within these efforts, I identify the engineers' conception of drywall finishing work as a layered practice with steps of craft well suited to human and steps well suited to machines, and their conception of the work embracing its contingency on an interplay between action, tool, and material. I discuss how these conceptions governed the system development to become a productive swing between the steps, quality, and technology, in which engineers shifted from unrealistic, full automation ambitions towards searching for a step well suited to a machine, from emulating worker skills with a machine towards focusing on the finishing material and quality, and from chasing advanced levels of autonomy, sensing, or motion towards experimenting with the material process, which enabled Basil to deliver industry-standard quality reliably. I identify that the finishers' explicit knowledge of the steps and tacit knowledge of finishing quality and evaluation became a central resource that informed the machine's initial development.

Next, I describe the observation of a present-day performance test called "1K," and discuss how Basil reshape part of drywall finishing into an iterative worker–machine sequence mirroring the system's sense-think-act architecture, and the conventional workflow into work-sharing between finishers and the machine, at the job level by reorganizing part of the division of labor, and at the wall level via a continuous feedback loop.[2] The final finishing quality, as I reveal, appears to be foremost

1 Introduced by scientist and philosopher Michael Polanyi (1966a) in the context of scientific knowing, "tacit knowledge" is a form of knowledge that "cannot be put into words" as opposed to "explicit knowledge" that can more easily be verbalized, codified, and taught to others. For more on "tacit knowledge" and "explicit knowledge," see Polanyi's (1958, 1966b) books *Personal Knowledge: Towards a Post-Critical Philosophy* and *The Tacit Dimension*. In this chapter, I use "tacit knowledge" as the knowledge or skills that drywall finishers acquire through experience but cannot fully verbalize how they practice or cannot easily teach to others. An example, as I further discuss in the chapter, is evaluating finishing quality. Whereas I use "explicit knowledge" as codified and easily transmittable knowledge in the trade, such as steps of drywall finishing.

2 Originated in the 1960s, sense-think-act is a robot system architecture that includes a sensing system that captures and converts inputs from the environment into a world model, a planner that determines a plan of action towards a goal based on this model, and an execution component that makes robot execute the plan, based on David Kortenkamp and coauthors' (2016) chapter. See more on sense-think-act and other system architectures in Kortenkamp and coauthors (2016).

in evaluating the success of this new worker–machine labor and facilitating the machine's acceptance in the industry. I show that the finishers' tacit knowledge of quality, along with their trained eyes and hands, played a critical role for both. Relatedly, I identify the emergence of finishers' tacit knowledge once again as a resource, this time informing the evaluation and further improvement of the system's performance. Both the early-state and late-stage developments show that, to put robotics in service of drywall finishing, drywall finishing knowledge is positioned in service of robotics at Canvas.

5.2 An Ethnographic Study of Robotic System Development: Part 1

5.2.1 Canvas

A thirty-five-minute train ride from Berkeley brought me to San Francisco's Mission District. I got off at the Sixteenth St. Mission station located under the Sixteenth and Mission Street intersection, offering passengers two steep escalator routes to the station's northeast and southwest exits. Each one leads to a small square, which exhibits the district's few characteristic elements at a single glance, such as slender palm trees and colorful murals on wall corners and demographics mixing the working class and the homeless. I stepped out from the southwest exit into San Francisco's fresh, cooler-than-normal summer air. After walking by a series of old residential buildings with shops interspersed with new condominiums, I arrived at the Canvas headquarters. The headquarters is a former industrial warehouse, one in a gallery of buildings spread between the eastern part of the district and the famous San Francisco Bay. The warehouse, as I will call it throughout this chapter, offers substantial floor space for the office, testing, fabrication, and storage needs of Canvas, which was already a well-developed robotics practice at the time of my study. I crossed the parking lot towards a large loading bay, checked in at the reception desk set up for the event I came to attend.

The event, held at the warehouse's loading-storage area, was the launch of Canvas's new robotic drywall finishing system called the Level 4 Targeted Spray (L4TS). The scene encapsulated the key elements of Canvas's technology design practice: the company's drywall finishing robot, Basil, live demonstrations of L4TS, which equipped the robot with image-based drywall joint detection and targeted spraying capability, employees that comprise engineers, managers, and drywall finishers who collaborated to develop these technologies, and various stakeholders who are interested in seeing or leasing them for jobs. Approximately one-third of the storage area was redecorated for reception, where two Basil machines demonstrated a few generic motions while the guests were joyfully interacting and celebrating. Next door to the reception, half of the wet testing area was arranged for the live demonstrations. At certain intervals, the crowd was invited into this space to spectate a third Basil machine and a drywall finisher from the Canvas team perform targeted spraying on a one-to-one-scale drywalled frame. The stakeholders with whom the company engages as their customers or users, included contractors and union workers, along with other curious tech professionals, who were not difficult for me to recognize and categorize in my head based on their self-presentation. During a brief pause between the demos, I finally shook hands with my two site contacts, who gave me access to study their robot and work.

A couple of days later, I returned to the warehouse to kick off the study. I spent the day orienting myself in the workplace, where I would be an insider studying Basil and the team of engineers, managers, and drywall finishers that developed it for the next eleven weeks. My site contact showed me the pantry, and then the office floor upstairs to get me settled. At the time of the study, Canvas employed a hybrid work model, allowing the team members to balance their office and home office time depending on their responsibilities. That particular day immediately followed the team's long,

tiring rush towards the L4TS launch. Therefore, the office floor was unusually quiet, not representative of what I would experience in the coming weeks. We sat, and I recapped the research plan I previously communicated to her, and shared my needs and potential key challenges up front. She suggested I start with three general meetings to get a sense of what was going on and emailed me the online meeting invitations. After she went on with her day, I settled at my desk in the office, and continued to tour the warehouse to familiarize myself with the dry testing, fabrication, and construction storage areas (Figure 1). The next day, I was formally introduced to the team along with my research agenda via a general meeting and an email.

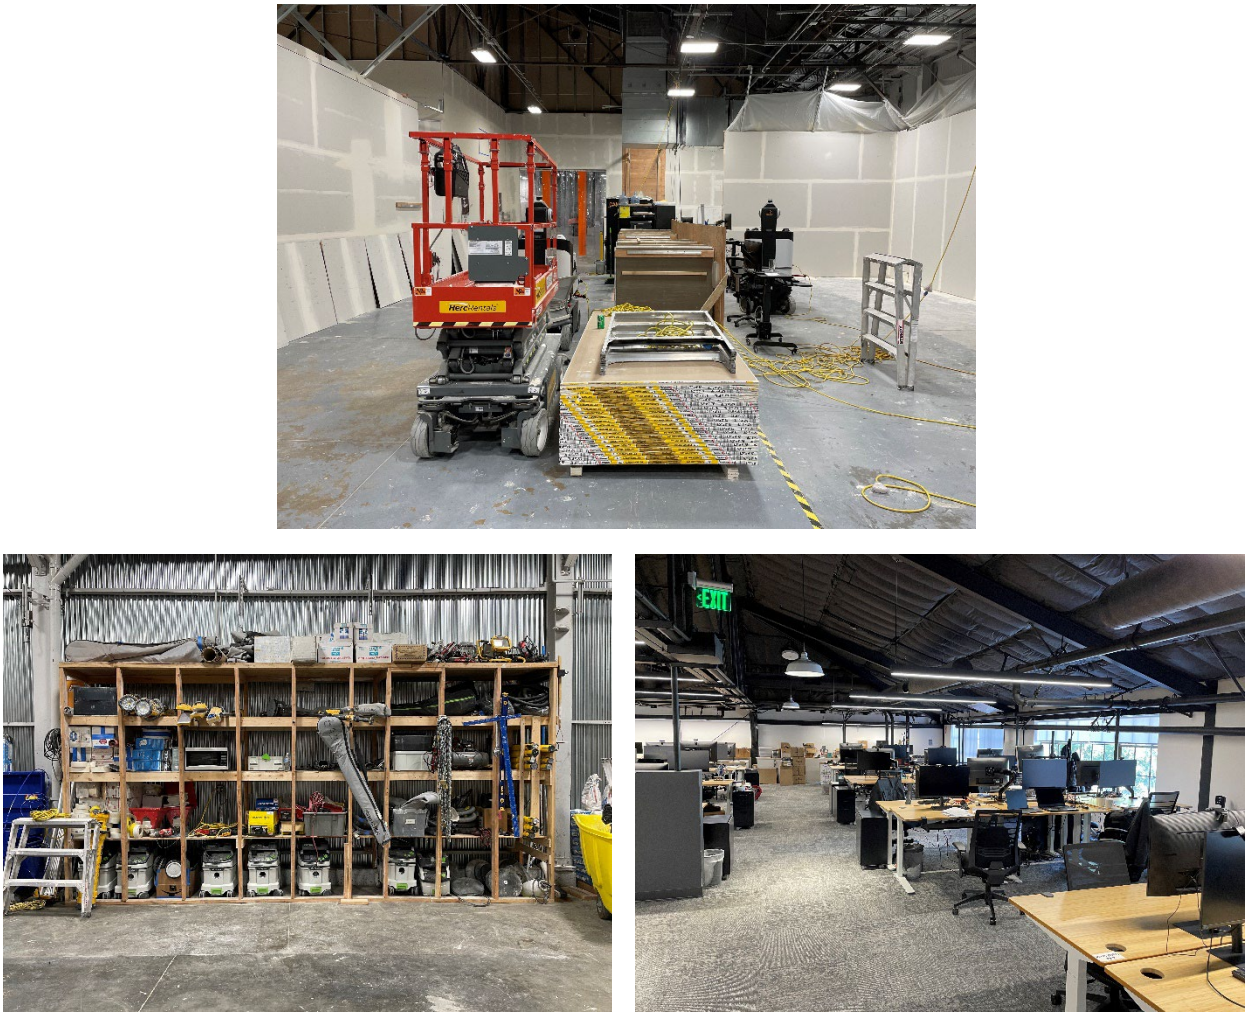


Figure 1. The wet testing area (top), the storage area (bottom left), and the office floor (bottom right)

5.2.2 Knowledge at the Core: Drywall Finishing

Work is a source of knowledge in the making of robots, which engineers and researchers seek to understand for developing and informing robotic systems as I pointed out in Chapter 1. This is also the case at Canvas, which I noticed during a few exploratory walks in the warehouse. Alongside engineers, Canvas employs a small team of drywall finishers, often referred to as “the crew,” who bring their trade expertise next to robotics.

On any workday, these professionals can be seen building mockup frames, hanging drywalls, or preparing material supplies on a typical six am to two pm construction schedule at the wet testing area, a job site quite different from the ones they are accustomed to. Part of the crew are also certified Basil operators who run the robot in robotic drywall finishing tests or jobs. These members, especially the lead of the crew, Drywall Finisher 4, provide the engineers with feedback about the system's quality of delivery during and after these events. Occasionally, Drywall Finisher 4 sits with the engineers in meetings where tests and the quality are discussed. The storage area, where the crew is based, resembles a giant hardware exhibition displaying the whole gamut of drywall finishing tools and supplies stuffed in gang boxes, carts, and shelves (Figure 1). My plan from early on was to immerse myself in the world that the crew built and learn about drywall finishing, which is at the core of the whole Basil project alongside the machine. In this endeavor, some crew members naturally became my teachers. Interviews with these members and their demonstrations of conventional drywall finishing were my invaluable learning resources.

Based on these sources, I describe drywall finishing, including the function of the task as well as the steps, tools, levels, and typical division of labor when it is performed on job site in the next four sections. By analyzing the steps, I draw out the skills that finishers develop and use in each step, and discuss these under two analytical categories: explicit and tacit knowledge. The knowledge of finishing steps, levels, and skills I unpack in these sections is the key material that the engineers at Canvas worked with during their early design efforts that originated Basil and the typical workflow with the machine, which I later describe and analyze in this chapter.

What Is Drywall Finishing, and Why Is It Needed?

To understand the function of drywall finishing, the task should be seen as a step in the sequence of other tasks, which produce a full interior partition together. The sequence goes as "framing, drywall installation, drywall finishing, and painting," as Drywall Finisher 4 listed to start our conversation. As the most seasoned finisher at Canvas with decades of experience under her belt, she was my teacher of the trade for an hour. Prior to our conversation, I had learned about framing and drywall installation from my time at the New Scaife Hall I described in Chapter 3. In a nutshell, while framing consists of constructing the basic structure of partitions with tracks and studs, drywall installation entails cladding this structure with fire-rated drywall panels. These panels come in four-by-eight-, four-by-ten-, or four-by-twelve-foot sizes, and, as I described in the New Scaife Hall workflow, workers cut them into custom pieces and hang them on studs in a composition that fits the shape of a partition. They use the marks on panels laid out in typical stud spans—every sixteen, twenty-four, or thirty-two inches—as a guide while fastening them to studs.

When two drywall panels connect together on a stud, they create joints, rough corners, and exposed nail spots, all of which need to be coated. Based on Drywall Finisher 4's description, drywall joints can be three types. As I checked myself later on a partition, when panels connect along their longitudinal edges, they form what finishers call a "factory seam" or a "field joint" (Figure 2). When two transverse edges come together, they make a "butt joint" (Figure 2). The long edges of panels are manufactured as slightly beveled inwards, which gives a factory seam a V-shaped profile. The short edges are slightly beveled outwards, which form a butt joint's reverse-V-shaped profile. Fire codes, building type, and wall height often determine whether the panels will be installed all vertically or horizontally, that is, in a "stand-up" or a "lay-down" layout. Occasionally, these two layouts have to be mixed, which causes a long and a short edge to connect and form a "bastard joint." When they meet in angles, they create inside and outside corners. Besides, panels have nails or screws on



Figure 2. A factory seam (left) and a butt joint (right) on a lay-down layout

them. Drywall finishing is the procedure to cover these joints and corners, along with the exposed nails or screws on panels with layers of joint compound, also known as “mud” in the construction industry. The aim is to minimize the uneven levels on partitions. Then the partitions are “hand[ed] over to painters” as Drywall Finisher 4 described.

Drywall Finishing Workflow and Tools

Drywall finishing is a “repetitious” task which involves performing multiple coating steps with a variety of tools over days to reach an almost leveled drywall based on Drywall Finisher 4’s description. A conventional workflow starts with mud and wall preparation, followed by taping, coating fasteners (also referred to as screw-on or nail spotting), corner coating, box coating; and ends with sanding. Later on, Drywall Finisher 1, a journeyman finisher in the crew, offered me a mockup demonstration of these steps, which I describe in this section. He began by introducing the hand and power tools associated with the task, which I introduce as I describe each step (Figure 3).[3]

Drywall finishing starts with mud and wall preparation followed by the taping step. As Drywall Finisher 4 described, finishers add water to a substance called the “joint compound” and blend the mixture until it reaches the right viscosity. They blend using a tool called a “battery-powered rotary drill” with a steel mixing pedal, as Drywall Finisher 1 pointed out. In parallel to mud preparation, finishers often perform a pre-inspection of issues on drywalls, such as fractures due to manufacturing faults or bulges linked to their installation, to ensure that the surfaces on which to apply mud are smooth. They conceal these fractures or bulges by adding an initial fill of mud with the “pan and knife,” which many regard as the two emblematic tools of the trade. After the mud and wall preparations are ready, the coating begins with taping. To show me how taping is done, Drywall Finisher 1 picked up a mechanical device called the “taper” or “bazooka” and stood in front of the partition. With his right arm at the bottom end of this long machine and his left arm at the front and under it, he pressed the tip of the machine against a butt joint and dragged it along to tape a short bit (Figure 4). The machine dispensed the first layer of mud with a strip of tape on top. He then said,

3 Based on the interview with General Manager 1, despite the tools specialized for every finishing step, pan and knife are still the dominant tools, used to execute many of the steps within small-scale residential construction. Whereas the specialized tools are preferred for faster execution and productivity in large-scale commercial jobs.



Figure 3. Part of the drywall finishing tool set

“you have to wipe down to make it [the tape] settle nicely without any faults or air bubbles,” while demonstrating how to do so with the six-inch-wide knife and the “wiping knife” for tapes on higher elevations (Figure 4). He added that this first layer should dry for about twenty-four hours.

Nail spotting and corner-coating often run parallel to taping as I learned from Drywall Finisher 4. Nail spotting entails coating the nails or screws exposed on drywall surfaces, and I observed how it is performed during Drywall Finisher 1’s demo drill. Drywall Finisher 1 explained that this coat is often executed in three layers via a specialized tool called the “nail spotter.” After he put the tool’s box-shaped head on a line of screws, he told me that the first coat is applied with the two-inch version of the tool and left to dry for a period. As he ran the head on a series of screws as if he were coating them, he continued and stated that the second and third coats, which use less mud, are added with the three-inch spotter and left to dry for a shorter period (Figure 4). To smoothen the edges, he followed with a quick brushing of the area with a sanding tool referred to as the “black widow.” (Figure 4). The other concurrent step is coating corners. As Drywall Finisher 4 explained, this step

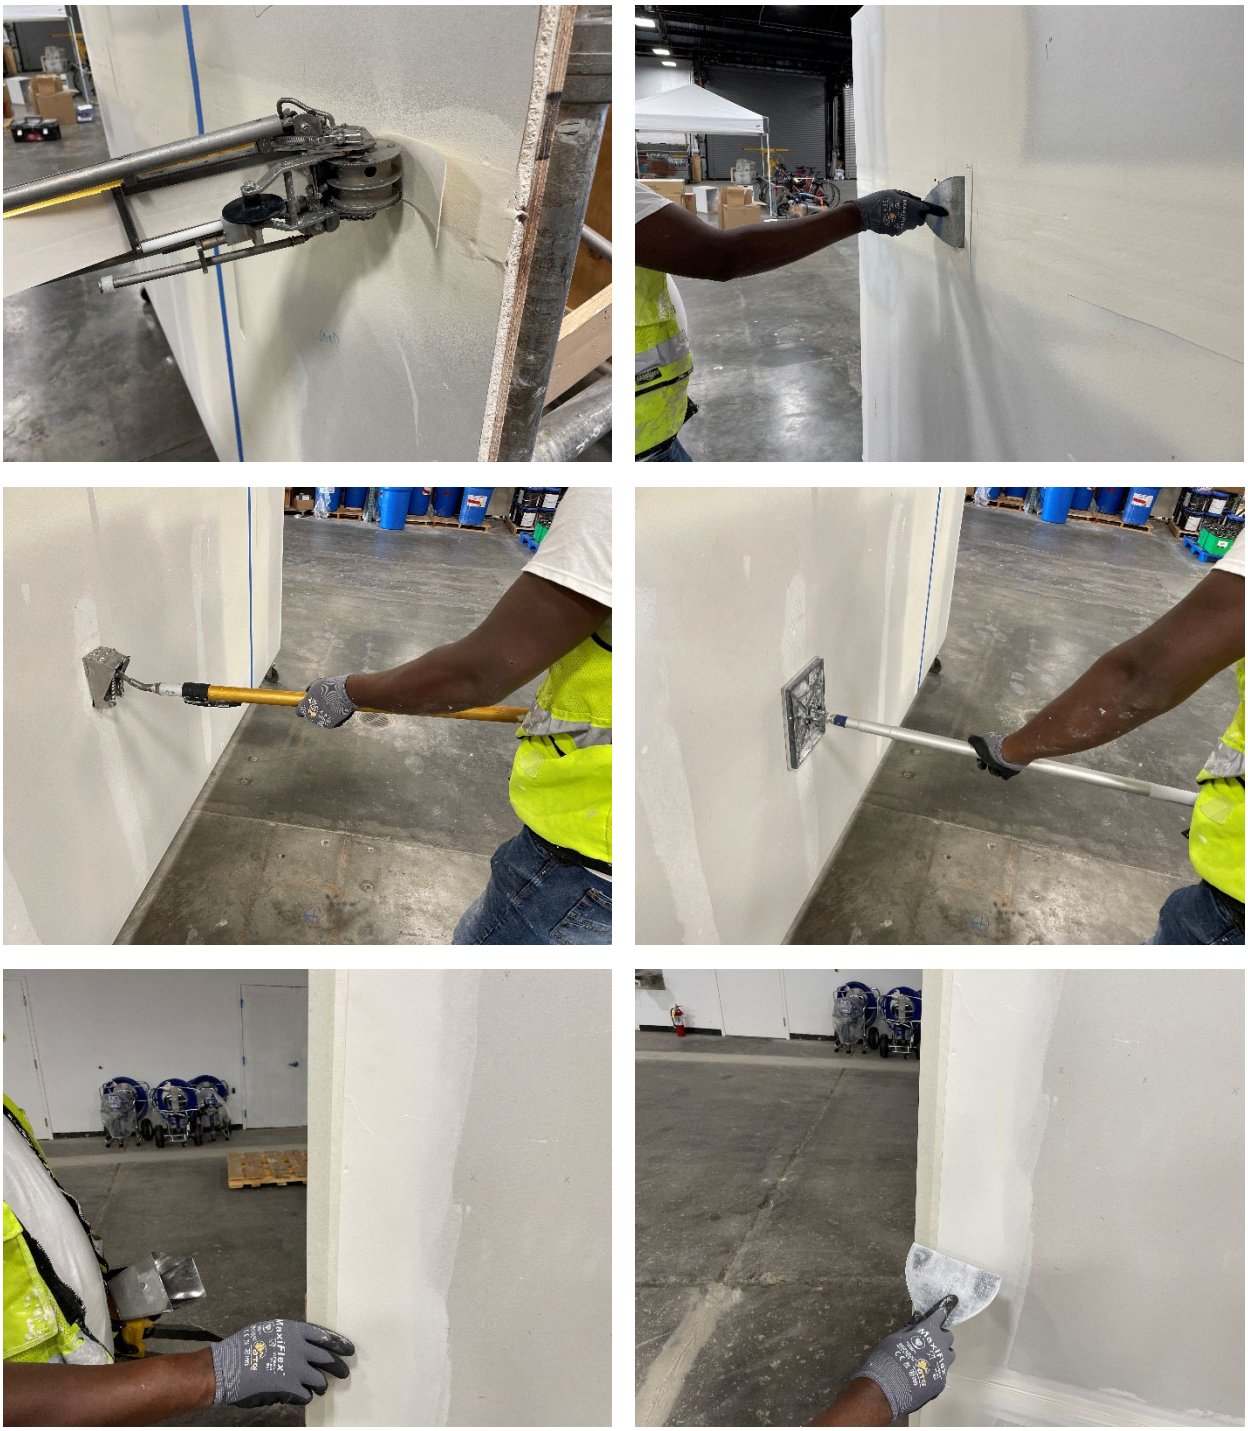


Figure 4. Taping and wiping (top row), nail spotting and brushing (middle row), and fixing the corner bead and wiping (bottom row)

consists in installing pre-formed frames called “corner beads” or tapes to define seamless corner profiles. To define outside corners, a paper corner bead with built-in tape is manually fixed on a layer of mud, which Drywall Finisher 1 demonstrated without mud (Figure 4). For the bead to sit fully, it is then coated with an additional layer of mud via the six-inch or wiping knife, as he showed me (Figure 4). For inside corners, tape is applied with the bazooka and wiped down with a special tool, the

“roller.” To make it evenly flat, the excess mud on top is removed by a tool called the “glazer.” Inside and outside corner coats should also dry for approximately twenty-four hours. Completion of taping, nail spotting, and corners marks the end of the first major layer of coating.

Taping, nail spotting, and corners are the baseline work for preparing the partition for the second major layer, box coating. As Drywall Finisher 4 explained, the purpose of box coating is to level the beveled profile of taped seams and joints as well as the taped or beaded corners with the rest of the drywall surface. It is done by adding three layers of coat, each applied with a specialized tool called the “box,” which comes in different sizes and dispenses mud in different widths and thicknesses. With the seven-inch-wide box in his hand, Drywall Finisher 1 explained that the first coat on factory seams is narrower and thicker, and he softly ran the tool along the center of the tape (Figure 5). This first coat serves to fill in the seam’s inward bevel. With the butt joint, he performed the same move on each side of the tape (Figure 5). This applies a slightly thicker mud to fill in the edges of the joint’s outward bevel. A second coat with the ten-inch box and a third coat with the twelve-inch one follows the first one he described, with twenty-four hours of drying in between. Coating is also applied in the same order to corners with a specialized tool called the “angle box,” along with smaller knives (Figure 5). If required, the rest of the drywall can be skim-coated with the ten- or twelve-inch knife, roll brush, or airless sprayer. Importantly, each coat is followed by a cross-off with the knife, which blends the mud where the seams and joints meet. “Every time you miss it [cross-off],” said Drywall



Figure 5. Box-coating seams and joints (top row) and corners (bottom row)

Finisher 4, the result is “less mud on the wall [in] some areas than [others].” The whole coating is finished after box- or skim-coating.

Sanding is the final step to get the partitions ready. “Relatively flat, but more so, you want a smooth wall,” was Drywall Finisher 4’s summary of what the finishing workflow should accomplish. Sanding is the final grating of coats to achieve consistent thickness so that the drywall fits her description. It may concern exclusively the seams, joints, and corners if only these parts are box-coated or more if the whole drywall is skim-coated. It is performed on all surfaces touched but especially the ones that ended up with more mud and appear as high spots by rubbing them with specialized tools, explained Drywall Finisher 1. He showed me the handheld sanding tools, including the “sanding poles” for surfaces, seams, or joints, the “sponge” for corners where the poles do not fit, and a power tool called the “vacuum sander.” The heads of these tools are fitted with sanding paper with appropriate grit depending on the amount of mud to remove. After sanding, all partitions must be checked for “any detail missed,” “any angles that look bad,” or “any nail spots that didn’t get hit [covered],” as Drywall Finisher 4 listed. This requires adding more mud with the knife, in a process called the “manual touch-up.” Any spot that is manually touched up is marked with a pencil or colored mud and left to dry. These marked spots are revisited and sanded the next day. After this final partial sanding, the partitions are ready for painting or other forms of decoration.

Levels of Drywall Finishing

Drywall finishing can be executed in five levels depending on the final decoration that will come on top of it.[4] Based on the sessions with finishers and a manual by the Gypsum Association (2015), a non-profit entity that provides technical information about drywall panels, the basic principle is that the level increases with the number of coating layers (Figure 6).[5] For instance, a partition with level 0

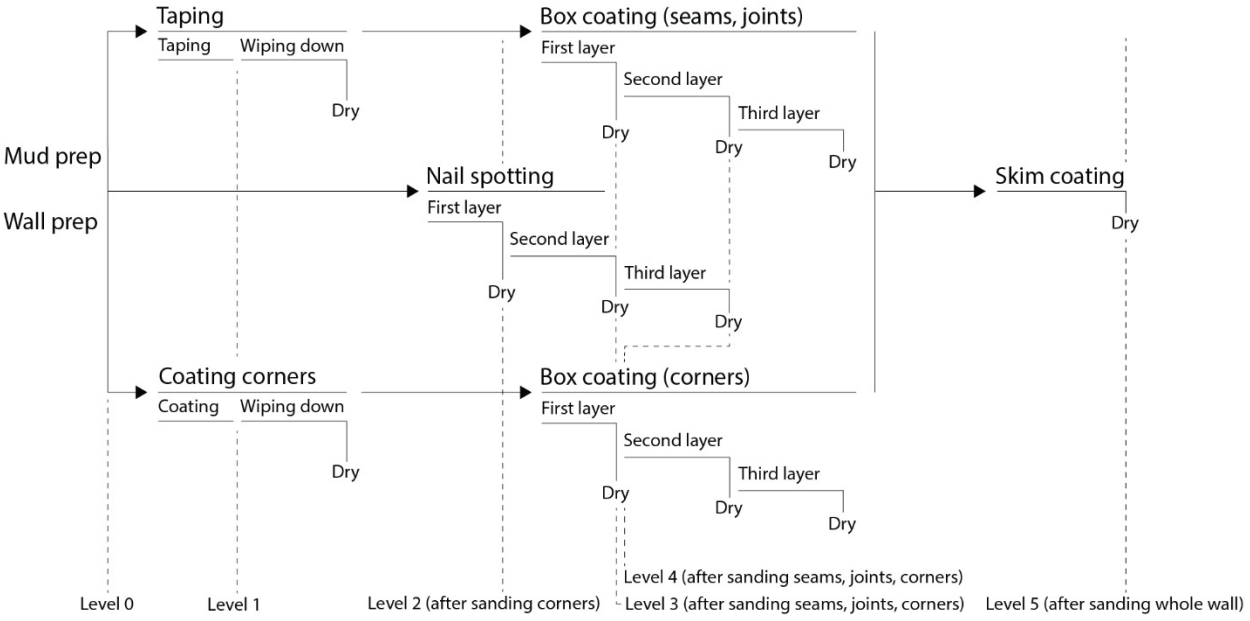


Figure 6. Steps and levels of conventional drywall finishing

4 This and the following details on levels are based on my interview with Finisher 4 and the manual mentioned in this section.
5 This manual is coded “GA-214-2015” and titled “Recommended Levels of Finish for Gypsum Board, Glass Mat, and Fiber-Reinforced Gypsum Panels.” The new version of the manual by the Association (2021) is coded “GA-214-2021,” and can be found on its website under “Publications & Documents.”

finishing is simply a bare drywall without coating. After taping and corner profiles are completed without wiping down the air bubbles or excess mud, the partition's finishing reaches level 1. When taping, the first layer of nail spotting, and the corners are done and all has been wiped down, the partition is at level 2, which is adequate for tiling on top. The first layer of box coat over the seams, joints, and corners and the second layer of nail spotting produce level 3 finishing, suitable for producing "heavy texture" partitions with paint or wallpaper on top. The second layer of box coat over the seams and joints, the first layer of box coat over the corners, and the third layer of nail spotting amount to level 4. This level of finish is suitable for creating "light texture" partitions with "flat," "light-tone" paint or wallpaper on top. This combination reduces an undesired effect known as "joint photographing," where the joints underneath coats are still visible. Lastly, an optional third box coat on the seams and joints of a level 4 partition and a skim coat over its whole surface makes it level 5, which enables the application of "non-flat," "dark-tone" paint or "other glossy decorative finishes" for a smooth wall with minimal joint photographing. All five levels can be encountered in buildings.

Division of Labor and Skills

Drywall finishing is a collaborative practice shared between workers, who have different levels of specialized skills for each step of the task, complementing each other in a concurrent or sequential order. Drywall Finisher 4 informed me that a commercial job is typically shared between five to six finishers depending on the project's scale. For a level 4 or level 5 delivery, three of these workers perform taping, nail spotting, and corner coating of the partitions' low parts, respectively, while the other two or three work on scissor lifts or scaffolding, each handling the same three steps on high parts. After these layers dry, the same five or six finishers share the box coating of seams, joints, and corners and sanding. As Drywall Finisher 4 noted, these different steps are delegated based on finishers' levels of experience or specialized skills. For instance, taping and nail spotting are more suitable for apprentices, who are beginners in the trade. This is because these steps, in which workers use the bazooka and nail spotter, are more tool-aided than dependent on freehand pan and knife skills. In contrast, actions that solely rely on hands, such as manual touch-ups or cross-off during box coating, demand dexterity with knives and higher levels of quality evaluation skills, which more experienced journeymen are expected to possess.

I distinguish these skills under two broad categories based on Drywall Finishers 4 and 1 and General Manager 1, who spent time in the construction industry, provided: knowledge of steps and tools, which involves more explicit forms of knowing what to do, and sensing of tools and outcome, which mobilizes more tacit forms of knowing or sensing how to do and whether the outcome is right.

Oftentimes, a combination of knowing and sensing are at play in finishing. For instance, as General Manager 1 explained, foremen have the knowledge of all steps, and are able to fit them into the trades' concurrent flow on a construction floor in a way that is efficient and safe. A combination of knowing the steps and certain basic rules, such as allowing a safe distance between their crew and others, together with sensing the ordering of equipment, materials, and tools that their crew's sequence will create on a floor enables these professionals to do this. Similarly, Drywall Finisher 4 mentioned mud preparation skills, which are common across all experience levels. Every trained finisher knows what to do on paper and what tools to use to prepare the mud properly and consistently, as well as how to do it in practice. They employ their knowledge of the mix—the knowledge of how to use the rotary drill and mixing pedals, as well as sensing whether the mud is at the right viscosity. Preparation primarily relies on procedural knowledge, including the right water-joint compound ratio or the speed, movement, and duration of mixing. In contrast, finding the

optimum combination of speed and length and determining the right viscosity is tacit knowledge, which Drywall Finishers 4 and 1 know by doing but struggle to explain to me.

A similar coupling of knowing and sensing is embedded in all coating steps and sanding. For finishers who perform taping, leaving the right amount of mud is linked to knowing when to apply the right pressure with the bazooka and when to pull it off the partition. This is knowledge of the step and tool. In practice, it is sensing or "listening to the machine [bazooka]," as Drywall Finisher 4 calls it, that helps finishers notice the blisters that the tool may cause when the tape does not sit on a sufficient amount of mud. For finishers who do nail spotting or box coating, knowing which size and settings to use for the nail spotter or the box, together with feeling the pressure while running the tool across the partition, is important to achieve the required width and thickness. Finishers working on these steps also need knife skills to perform any manual adjustments, wiping, or cross-off. "It seems easy, but it is not easy how to hold that knife," said Drywall Finisher 1 to emphasize the intricate nature of using this tool. After selecting the appropriate size, the proper way is to put the index finger at the center and use it, along with the palm and wrist, to control the pressure that determines thickness, as he demonstrated (Figure 7). The tool is navigated by shifting and pressing down the finger in the intended direction (Figure 7). For its part, corner coating is about knowing the basic moves with the roller, developing "an eye for straightness [of the bead]," as Drywall Finisher 4 mentioned, and, again, applying the right pressure to make the bead sit nicely. Finally, finishers working with the sanding pole or vacuum sander rely on a combination of knowing which sanding paper to use and having good control of the speed and pressure to remove the right amount of mud without scraping the drywall.

Parts of the task that a finisher's sensing becomes most critical is the evaluations of the drywalls before finishing and of finishing after each coating step and after sanding as my conversations with General Manager 1 and Drywall Finisher 4 revealed. Foremen and some journeymen use quick eyeballing skills during pre-inspection to spot "edge cases," as General Manager 1 described. An edge case could be a "bad wall" with fractures or a "bad joint" with bulges that requires a manual touch-up to make the wall ready for finishing. After each step, foremen and journeymen combine visual inspection with touch to spot "any mishaps" in critical areas, as Drywall Finisher 4 added. Drywall Finisher 4 listed these critical areas as "blisters [on tapes]," "[uncovered] nail spots that did



Figure 7. Drywall Finisher 1 demonstrating how to use the knife

not get hit," "angles [inside corners] that look bad," "[where] the boxes [box coats on seams and joints] meet together," and "low parts and fractures [on the drywall]." Eventually, all these areas should be "relatively flat," as she put it. Experienced finishers usually inspect blisters and uncovered nail spots visually or by gently caressing the surface with their fingers. To check mud thickness, the consistency of thickness, and texture after coating the seams, joints, and corners and after sanding during post-inspection, they use the same techniques and determine whether any area needs more mud or sanding before handing the partitions over. Fascinatingly, the workers can see and feel the quality of the finishing. This specific tacit knowledge has an instrumental role in informing Basil's early-stage development and in evaluating the machine's current performance during the 1K test, which I discuss in the next two sections respectively.

5.2.3 From Conventional to Robotic Drywall Finishing Workflow: Basil

A Brief Description of Basil

The current Basil is the latest in a line of drywall finishing machines, including its earlier versions and a couple of early prototypes, all developed by the engineering team during the company's initial years (Figure 8). Except for a few minor changes over time, the following assemblage has been the software-hardware skeleton of this semi-autonomous machine. The machine's main body is made of a human lift repurposed into a mobile base with 2D and 3D LiDARs, RGB-D cameras, and a height sensor. A collaborative robot arm, which works with an airless sprayer (which sprays the mud on partitions) or a vacuum sander (which smoothens the sprayed parts) is mounted on this base. The base moves the arm on a construction floor and lifts it vertically to different heights on a partition. Other than the arm, the base also carries a compressor for the sprayer and sander, the vacuum part of the sander, and the system's onboard computer. The airless sprayer and a mud drum (a large-scale bucket that supplies the joint compound to the sprayer) are external to the base, and both items are carried alongside the base on a custom-built wheeled cart. A controller consisting of an off-the-shelf tablet and a gamepad for monitoring and controls is the last major hardware. A software stack that handles the communication between the base, sensors, and software, communication and control between the different modules of the software, localization, planning of the task, seam detection, and planning and execution of the arm sequences is deployed onto this customized vehicle. Attached to this deck is a user interface for giving inputs to the system and receiving outputs from it.

With this assemblage, a finisher can execute level 5 finishing with robotic spraying, which, as I explained, is the appropriate base for applying a non-flat, dark-tone, or glossy paint to deliver smooth partitions in commercial and, more often, healthcare buildings. The steps of this workflow include conventional taping, nail spotting, and a layer of box coating, followed by robotic spraying and sanding. The newly launched L4TS system that the team celebrated and on which I focus in this chapter extended the machine's capability to deliver level 4 in addition to level 5, mainly via additions to its software. The system eliminates box-coating; however, the workflow includes manual sanding performed with a vacuum sander.[6] L4TS was a response to the industry's preference for level 4, which can also help achieve smooth, light texture partitions with a slightly lower quality than level 5 finishing but in less time and with less labor and materials. This strategic move aimed at increasing the machine's appeal for the industry.

6 At the time of this study, Basil could execute sanding on entire wall for Level 5 finishing, however, the system did not yet have the capability of sanding the seams only for Level 4.



Figure 8. Basil (bottom) and a precedent machine (top)

A Productive Swing Between Steps, Quality, and Technology

In this section, I trace and describe the two main lines of effort within the early-stage development of Basil in the past drawing from my conversations with the engineers. The first one was the engineers' learning and examination of conventional drywall finishing steps, which I outlined in the previous section, through the lens of human and machine capabilities. The second line of effort was their work on operationalizing the finishing quality when mud is sprayed and sanded with the machine. Within these lines of effort, I look into how these engineers conceptualize drywall finishing work based on their learning and examination, and how these conceptions come into play in Basil's initial development. Further, I identify the emergence of finishers' knowledge of steps, quality, and quality evaluation informing these lines of effort.

As my conversations with the engineers revealed, a line of effort that shaped the machine's technical underpinnings is learning conventional drywall finishing and examining its steps based on a distinction between human and machine capabilities. "People and robots actually share very few

commonalities," said General Manager 3 as she unpacked the way of thinking underlying this practice. She continued by summarizing the key distinction: "people are good at adaptability… versus robots are good at repeatability, and doing the exact same thing very precisely." This broader distinction guided the engineers in the early days of Canvas while studying the conventional finishing steps. The knowledge came from Drywall Finisher 4. They assessed the steps in terms of what is difficult for humans and suited to machines at the same time. The notion that "everything that requires you [as a human] to touch the wall is hard," in Engineer 7's words, was the key inference derived from that search. "For [a] machine, applying constant pressure for hours and hours is not hard at all," she said, calling on this principle to explain why the engineers chose level 5 as the machine's initial primary function. The reason is that the large wall surface to spray and a large quantity of spray to sand amount to the perfect job cycle for a machine.

Looking at drywall finishing through this human/machine distinction illustrates the engineers' grounded conception of drywall finishing work, embracing the reality that some steps are well suited to human capabilities, such as dexterity, improvisation, and craft while others are better suited to those of machines, such as handling layers of coating and quantity. I observed that this guided the engineers to take a couple important departures from certain prevalent technology design approaches in robotics.

The first departure was from unrealistic, full automation ambitions towards searching for a part appropriately suited for a machine. Drywall Finisher 4's explicit knowledge of the coating steps helped to answer the engineers' key question here: what is possible with a machine in drywall finishing and how would it fit? The answer that came out of an initial ideation, prototyping, and experiments was a machine combining a base, robot arm, airless sprayer, and vacuum sander could come into a job after taping and nail spotting and help with the coating and sanding of large surfaces. This would reframe the conventional manual workflow I previously described in this chapter into a combined manual and semi-autonomous one and distributed its steps between a team of finishers and the machine, which complement each other's capabilities. In this distribution, the finishers would handle the mud preparation, wall prepping, taping, nail spotting, and coating corners, which demand more dexterity or improvisation (Figure 9). For its part, the machine would perform the full spraying and sanding for level 5, eliminating box coating and manual sanding, or targeted spraying for level 4, eliminating box coating (Figure 9). As final steps to complement the machine work, the finishers would also carry out final blend-ins, manual touch-ups, and quality evaluation, which require dexterity and qualitative inference (Figure 9). The knowledge of steps worked as a reference for engineers while positioning the machine and conceiving the software-hardware that enabled them to eliminate box-coating.

The second departure was from emulating worker skills with a system towards focusing on the mud and finishing quality, in which Drywall Finisher 4's knowledge of coating steps became key. In technical terms, the engineers chose not to go for a learning-based approach, which would involve data curation and training a model to emulate a skill. "You want to define a new process that doesn't mimic the human," is the principle behind this departure, as General Manager 4 explained, firmly asserting that "mak[ing] a robot that do[es] sense stuff as a human is never going to work." Therefore, instead of focusing on finisher skills, the engineers' starting point in designing the machine was the mud and the final quality of the finishing. The very first ideas for the machine, as General Manager 3 recalled, did not aim to "do it [drywall finishing] how a person did it" but to "achieve the result that a person achieved." This departure also led to innovating the conventional finishing steps prescribed by the Gypsum Association's manual. "There is, from the start, a departure from what manual work has been" said Engineer 7 as she explained that the hand tools

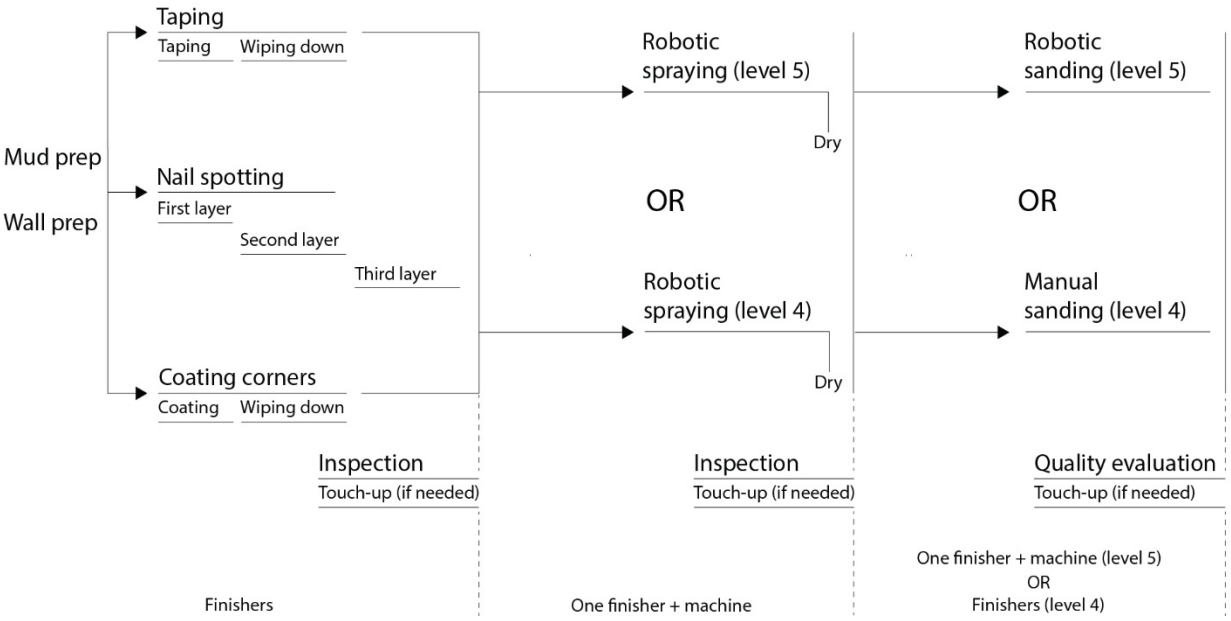


Figure 9. Steps of drywall finishing workflow with Basil for level 4 and 5

did not become the customized end effectors and the robotic workflow did not end up following the material process exactly. More specifically, the three box-coat layers applied on top of dry tape in conventional level 4 finishing became a single targeted spray layer on wet tape in the robotic L4TS workflow. As seen in this case, knowledge of steps enabled the engineers to redesign the coating sequence.

In light of these notions, the engineers built the first few units of Basil, and that was when the second line of effort focusing on material and operationalization of quality began. With some form of advanced prototype in hand, they could now spray mud and sand it in a relatively reliable way. So, the focus shifted to how to make the outcome perfect, specifically up to industry-standard quality, and the delivery efficient and reliable, so that it is within reasonable amount of time and consistent. That is when the team started a new effort, which Engineer 12 characterized as “focused on finish and working backward from that, and saying, ‘how can we best solve that [delivering this quality with the machine] in any way possible?’” Technically speaking, this reverse-engineering was about fine tuning the software and hardware based on studying the mud and evaluating the machine’s spraying and sanding results.

This outcome-driven approach to system development indicates the engineers’ conception of drywall finishing as a chiefly material-oriented practice contingent on an interplay between action, tool, and material. This conception guided the engineers towards shifting the importance from chasing advanced-level autonomy, sensing, or motion towards experimenting with the material process that would enable the machine to reliably deliver industry-standard quality.

The experiment involved an interesting modeling process, in which I found a fertile interplay between Drywall Finisher 4’s knowledge of the mud and finishing quality and Engineer 7’s knowledge of quantitative modeling and evaluation. “There is no textbook or norm, [all is] done working with [Drywall Finisher 4], … understanding how they [the finishers] evaluate things and then trying to figure tools of how to measure that,” said Engineer 7 while explaining the process. Essentially, the two professionals from engineering and construction translated the physical process that takes place

between the arm, software, mud, and drywall panel into a model with quantitative parameters related to the system and qualitative aspects related to the final effect on the partition. Engineer 7's job consisted in defining "input parameters" that could be controlled on the software and hardware side and whose manipulation would change the output parameters and specs.[7] Output parameters are related to the effect and the associated dimensions of the sprayed mud on the partition, such as mud thickness, the consistency of thickness, and fan width, which are measurable with basic construction tools like the tape measure and a thickness measurement tool called the "wet film comb."[8] These measurements, measurement categories, and tools to evaluate quality came from Drywall Finisher 4. Output specs are a list of quality properties concerning the result after sanding, which Drywall Finisher 4 has inspected and evaluated countless times on job sites visually and via tactile sensing, including mud thickness, the consistency of thickness without any raised or depressed spots, and fractures.[9]

To determine how these parameters and specs work in relation to each other, Engineer 7 and Drywall Finisher 4 collaborated in a series of experiments. These experiments involved iteratively testing the machine's spraying and sanding with specific values or ranges in a perfectly conditioned environment and evaluating the outcome "to tie down [these] output parameters to [the] specs that [Drywall Finisher 4] gives," as Engineer 7 described. "[Drywall Finisher 4] comes and evaluates the results of the experiment and helps me to calibrate [the input parameters]," she highlighted and helped me trace how qualitative evaluation is connected to modifications in the software. The model that resulted from this collaborative effort standardized the spraying and sanding processes and helped to control their outcome more effectively. It enabled the team to initially get the system to achieve a reliable industry-standard finishing quality and, later, to fine-tune it in an iterative cycle of tests, finishing quality evaluation, and software-hardware adjustments. From knowledge of quality and quality evaluation came the specs against which the machine's outcome was tested and benchmarked until it became reliable.

What we may see through these conceptions and processes is, developing a machine for construction is not confined within the boundaries of robotics but positioned at the intersection of engineering and construction at Canvas. It is not really taken as skill translation via learning-based techniques in robotics but rather a knowledge transaction between robotics and construction. Construction knowledge, in particular the parts that focus on the steps, skills, and quality are positioned as the central resource to inform the initial development of the machine. As we see in the early ideation, knowledge of steps and skills became references for engineers in searching a suitable part for a machine and conceiving software-hardware that would fit to perform it. As we see in the experiments by Engineer 7 and Finisher 4, knowledge of quality and quality evaluation became a technology design instrument helping to generate the model, which operationalized the effect that the machine produces on drywall.

Sense-Think-Act + Evaluate: Work-Sharing Between Finishers and the Machine

Early one morning, the wet testing area was filled with excitement and some nervousness. A team of engineers, along with Drywall Finisher 1, were finalizing the preparation for one of the last tests of L4TS, called "1K," before the system would be used on a job site.

7 This information is based on the interview with Engineer 7.
8 Ibid.
9 This information is based on the interviews with Engineer 7 and Finisher 4.

Based on the observation and detailed descriptions of this test, I discuss three important points regarding the distinct way of working with Basil in this section. The first one is that the machine's sense-think-act architecture turns part of finishing work into an iterative worker–machine sequence. The second point is that the machine's integration reshapes the conventional workflow I described earlier in this chapter into work-sharing between finishers and the machine. The last point is, similar to the machine's early-stage development, final finishing quality was a key element and evaluated by expert eyes and hands next to engineering measurements to determine the workflow's success. In this evaluation practice, I traced the emergence of finishers' knowledge of quality and quality evaluation, this time as a measurer of machine's performance.

The wet testing area is a large-scale design-of-experiment setting, having the mixed character of a relatively neat mockup of a construction site and a robot test zone. Mockup frames fitted with drywalls are the dominant elements in this space, dividing its bare concrete floor into zones. On a given day in these zones, in-house finishers can be seen on lifts hanging or prepping the drywall panels steps away from a group of software engineers operating a machine for testing or debugging. Some days, an apprentice from a drywall contractor may be receiving training in operating Basil or engineers may be demoing the system to curious potential customers. For the 1K test, the test team circumscribed a lane along one of the long straight partitions (Partition 1) with panels hung in a stand-up layout (Figure 10). This was part of engineering the ideal conditions, including a sufficiently long straight wall, an experienced operator, and no concurrent work, to test the machine's ultimate performance in these conditions. Other than Partition 1, the test would also be conducted on a short stand-up partition (Partition 2) and a long lay-down partition (Partition 3), both with turns. On each of these partitions, taping and nail spotting, which I identify as "pre-machine steps," were completed prior to the test.

This multi-day test would involve Drywall Finisher 1 running the targeted spray workflow as closely as he would do it in an actual job to eventually reach 1K square feet minimum target.[10] While the workflow would continue, the engineers would collect data to evaluate a set of performance measures, including the machine's reliability, the operator's comfort level with the UI, the productivity of the workflow, and the quality of the spray. In the next step, three in-house finishers, including Drywall Finisher 1, who would operate the machine, would manually execute the sanding with the vacuum sander. As the last step, Drywall Finisher 4 would go over the seams and joints to evaluate the final finishing quality.

Drywall Finisher 1 kicked off the test by starting the Basil machine. He entered the dimensions of Partition 1 on the UI tablet, positioned the base by using the UI gamepad, and pulled the wheeled airless sprayer cart closer to it.[11] Then, he fixed a few spots on the partition by manually adding more mud and readied the partition for the upcoming targeted spraying. Meanwhile, Engineer 1 was preparing her data collection tools, including a notebook and a shared spreadsheet document, on which she would be recording any faults, bugs, or unexpected software-hardware issues related to machine reliability (Figure 10). The other tool was the software the engineers used to monitor the system, through which she would check whether Drywall Finisher 1's UI interactions went smoothly.

10 1K square feet target was deliberately set higher than the standard daily amount on job sites. This was because the 1K test would be performed in ideal conditions, including an experienced operator, a more open floor, and no other concurrent work nearby. Targeting 1K was also part of Canvas' goal to improve the current standard.

11 The detailed description of Finisher 1's workflow, depicting his actions in both physical space as well as on the system's user interface, is based on a synchronous observation in-person and an asynchronous observation of his UI sequence through the software that Canvas uses to monitor all the machines, called Formant, developed by a company with the same name. For more information, see formant.io.

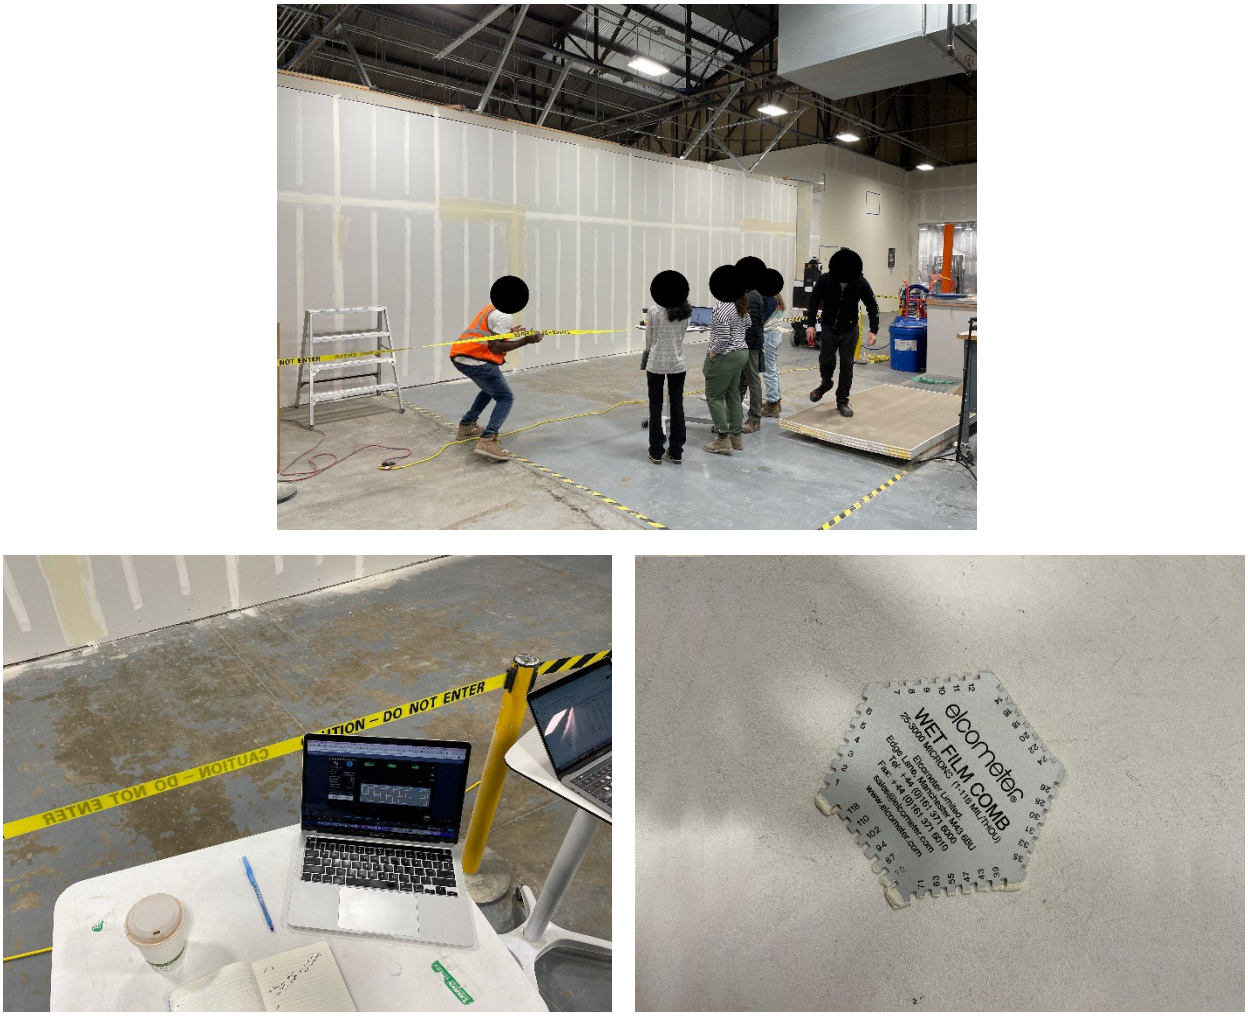


Figure 10. The team in front of Partition 1 (top), Engineer 1's test data collection tools (bottom left), and Engineer 3's wet film comb (bottom right)

Next to her, Engineer 2 was focused on the same shared spreadsheet, on which she would be keeping an activity-time log of the workflow for the productivity assessment. Engineer 3's job would be to take the thickness and fan-width measurements of the sprayed mud on joints between and after spraying. He stood with a tape measure and a wet film comb in his hands (Figure 10). I was set to observe and document the workflow through notes and photos. Drywall Finisher 1 had already mixed the mud and set the sprayer's pressure. With the UI tablet in his hands, he was ready to start.

"Planning" and "execution" are the two main steps of the spraying workflow.[12] In the planning step of the 1K test, Drywall Finisher 1 first set the "hazard height" to one hundred and forty-four inches—the maximum height the base should safely lift to in this case—and checked the "drywall install" option, which was set to "stand-up" by default. The machine's arm is programmed to work on either the base's right or back side. Partition 1, with its long straight profile, can be worked off the side all along, so the "work off the back" toggle was switched off. Along with the partition dimensions, this was all the input the system needed from Drywall Finisher 1 before he finger-pressed the "generate

12 Before the 1K test, I learned how to operate Basil in a training session that Canvas offers to finishers from outside contractors, and I learned more about the L4TS workflow in a test that inspects the main functions of a brand-new Basil unit.

plan” button. This prompt sent the input to the software module that handles the planning of the task. This module divides the partition into columns and columns into a non-uniform grid of cells called the “workspaces,” each labeled with the letter C and a number. Seams and joints falling into each column can be sprayed with the base parked in one appropriate position on the XY plane, and each workspace in that column can be sprayed with the base lifted to an appropriate Z coordinate. In this case, the module generated thirty-three workspaces in eleven columns, and the UI displayed these overlaid on a graphic representation of Partition 1 (Figure 11). Upon seeing this digital wall on the screen, Drywall Finisher 1 started the execution.

The execution step starts by selecting the workspaces in a column. Drywall Finisher 1 selected the C0 and C1 workspaces in the first column and pressed the “start execution” button. He skipped C2 as he saw that Partition 1 had no seams or joints at that height. Because the first column was the partition’s rightmost edge, he located an additional cardboard-like surface attached to a tall standalone post. I realized that this was meant to block the mud that would be sprayed outward. Next in the execution step is positioning the base. For this, Drywall Finisher 1 checked the view on the UI, which showed a rectangle representing the base on the Cartesian plane, and fixed it at a comfortable zoom distance before beginning to maneuver the base. As he was maneuvering, he looked carefully at the base in both physical and Cartesian space. For the arm’s movement to cover each workspace, the base should be positioned within a range of distance to the partition. The UI guided Drywall Finisher 1 in doing so with five reference lines indicating the proximity of the base to its periphery (Figure 11). Based on these lines, Drywall Finisher 1 quickly set the side distances. Then, he manually adjusted the wheeled sprayer cart, finalized the positioning, and pushed the “confirm position” button. The UI showed the “ready to start execution” window, which meant that the system was ready to detect the seams and joints within C0.

On the “ready to start execution” window, Finisher 1 pushed the “start” button. In response, the system’s RGB-D cameras captured an image of C0, and this image input was sent to the software’s seam detection module. The entire first column was smaller than usual, with only one butt joint to be sprayed. In the review window on the UI, Drywall Finisher 1 saw that the whole extent of that one joint was not in the image, which made him eliminate C0 and move on to C1. Accordingly, the system lifted the base to the appropriate Z coordinate, and the same actions were repeated in the same order. An image of C1 was captured, showing the full butt joint and overlaid with a yellow transparent box and a blue axis line on the review window (Figure 11). Drywall Finisher 1 carefully reviewed the screen, manually adjusted the axis to align with the joint, and pushed the “spray 1 joint” button to start the arm’s movement sequence and spraying. For butt joints, the arm performs two adjacent passes to make a wider coat, so it did two passes. After finishing the passes, the arm went back to its default position, and the UI notified Drywall Finisher 1 of the completion of C0, C1, and C2. The first column was successfully executed, and there were ten more to go. This worker–machine action series I described is the typical sequence that repeats for each column. After executing the first column, Drywall Finisher 1 selected C3, C4, and C5 in the second column, once again positioned the machine at an appropriate distance, and ran three rounds of seam detection and spraying to complete these workspaces. Aside from a few breaks for taking measurements and handling minor errors, the remainder of Partition 1 was finished that day by repeatedly running this sequence (Figure 12).

An important point that emerged as I watched this workflow concerns the way the system’s sense-think-act architecture frames part of drywall finishing. Through this functioning, the system recharacterizes part of drywall finishing into working iteratively with system-generated divisions, which is its essential way to sense and represent the partition. Each sequence I observed is



Figure 11. Drywall Finisher 1's UI sequence (top grid), Partition 1 columns and workspaces (left column, top), the base on the Cartesian plane (left column, middle), and butt joint within C1 captured by the system (left column, bottom)

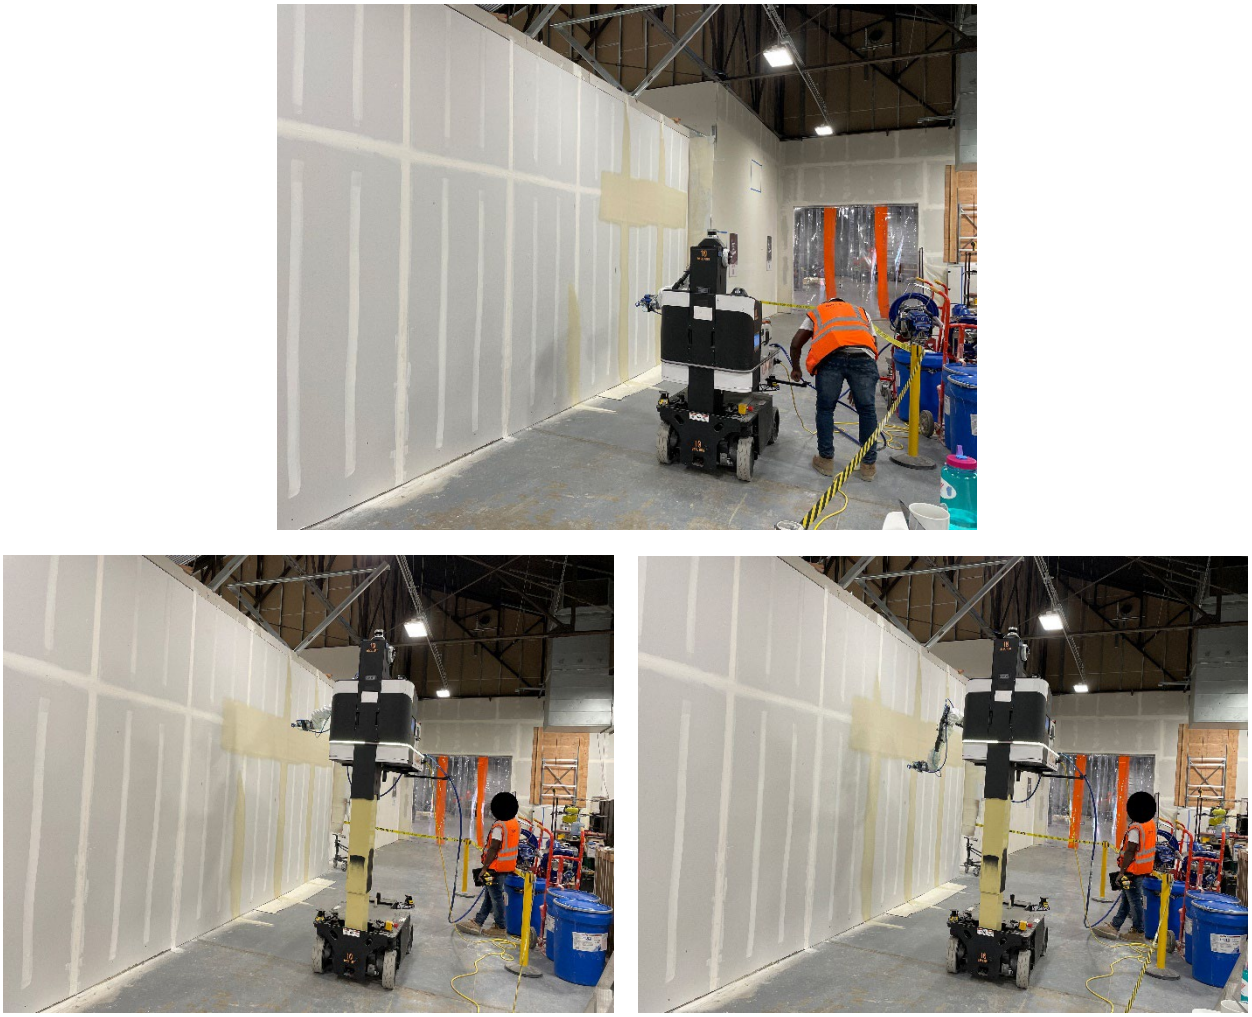


Figure 12. Basil spraying part of the factory seam in C9 (top), the butt joint in C10 (bottom left), and the rest of the factory seam in C10 (bottom right)

essentially made up of sense-think-act cycles on the system side. After generating the plan, the horizontal positioning of the base relative to a column is the “sensing” part of the cycle, while the detection of the seams and generation of the arm’s waypoints corresponds to the “thinking” aspect, and the execution of the arm’s motion and spraying constitutes the “acting” part. Each workspace is executed in one sense-think-act cycle, followed by the vertical self-positioning of the base relative to the next workspace, which is the “sense” part of a new cycle. After a column is finished, shifting the base to the next column represents the “sense” step of the new first cycle, and the sense-think-act cycle repeats in every workspace of this new column.

The other notable point I observed during the 1K test workflow is the sharing of work between finishers and the machine at various levels. When zoomed out to the job level, one can see that the robotic targeted spray is only one step, situated between the pre-machine steps (taping and nail spotting) and sanding, which comes after the machine (Figure 13). This three-part labor that brings a bare drywall panel to level 4 finish is conducted in order by finishers performing the pre-machine work, then a finisher using the machine, and finally finishers doing the post-machine sanding. When zoomed in to the wall level, this is facilitated through a continuous feedback loop between one finisher and the machine in sense-think-act cycles, in which the finisher and the machine give each

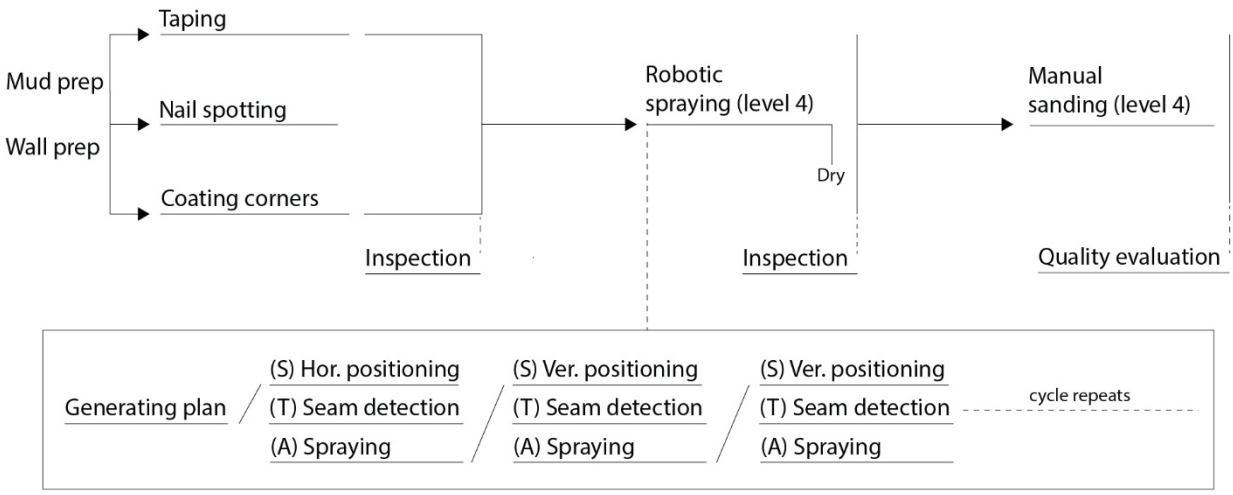


Figure 13. Level 4 drywall finishing workflow with Basil, including pre-machine steps, robotic targeted spray working in a feedback loop with the machine through sense (S), think (T), and act (A) cycles, and sanding

other cues via inputs and outputs (Figure 13). A good example is the horizontal positioning of the base: the finisher handles the positioning through inputs via the UI, and the machine provides positioning feedback to guide the finisher. In seam detection, the machine provides the detection output, and the finisher approves or adjusts. Thus, at the job level, the machine reorganizes part of the division of labor in drywall finishing. In addition, it translates the work at the wall level into a series of UI actions for the finisher to execute the sense-think-act cycle with the machine. As mentioned earlier in this chapter, this eliminates the box-coating step, and the UI comes in as a new tool replacing the boxes.

The last significant point I drew from the 1K test is the diligent effort spent on quality evaluation during and after spraying. This, in addition to Engineer 7 and Drywall Finisher 4's iterative cycle of experiments described previously in this chapter, was where I again noticed how central the in-house finishers' tacit knowledge of quality evaluation, which goes hand in hand with the engineering measurements, is to system development. Because Canvas essentially services project owners, general contractors, and drywall contractors, the ability to deliver an industry-standard quality, together with the acknowledgment of this ability by those parties are important benchmarks for Basil to be adopted. Therefore, the Canvas team continuously evaluates the machine's performance based on the quality of the outcome it produces. I observed during the 1K test that one way the team does this is through numbers. Before the test, the team outlined a list of success criteria in the form of the percentage of outcomes that meet the satisfactory range of alignment, fan width, and mud thickness measurements. As part of this evaluation, after spraying three columns on Partition 1, Engineers 1 and 3 took these measurements to assess a few sprayed seams and joints on the fly (Figure 14). At the end of each spray day, Engineer 3 measured these for all the sprayed seams and joints and penciled his comments on the partitions (Figure 14). The second, and perhaps a more important way I observed the team evaluating the outcome and, consequently, the machine's performance is through expert evaluation. How a measurement range is determined satisfactory is by carefully associating the numbers with what is qualitatively acceptable to the in-house finishers.

This evaluation primarily relies on tacit forms of knowing and sensing whether the mud on the partition is suitable as I previously highlighted in this chapter. On sites, this is commonly done via visual inspection by experienced tradespeople, as General Manager 1 explained: "[how you gauge quality] is still very much a, first you do a process and then you'd have a visual job site inspection to



Figure 14. (1) Engineer 3 taking measurements during and writing notes on a partition after spraying (top row), (2) Drywall Finisher 1 sanding a partition and the team performing the post-inspection after sanding (middle row), and (3) Drywall Finisher 4 writing comments on a seam and using touch for quality evaluation (bottom row)

confirm." In the 1K test, I saw Drywall Finisher 4 demonstrate one such inspection before spraying to inspect variables related to drywall panels that might affect the outcome and after both spraying and sanding to determine whether the final finishing was up to industry standard. During the pre-inspection, she checked each panel to spot manufacturing faults, such as tears, exposed screws, or

fractures on tape, that required more mud, accompanied by Engineer 1 taking notes. She also used tactile inspection as a quality evaluation technique on some parts within her reach, checking for the presence of any raised or depressed spots on the tapes. After spraying, the two ran the post-inspection to determine whether the mud thickness, the consistency of the thickness, and the texture were suitable and, accordingly, whether the sprayed seams and joints were ready for sanding.

The final step conducted on the next workday was sanding, which was another instance of the in-house finishers' tacit knowledge of quality and quality evaluation worked as a reality check to measurements and guided further improvement. This particular sanding process was an experiment in which Drywall Finisher 1 tested different vacuum sander speeds with different sanding paper grits on targeted spray (Figure 14). The aim was to find out which combination would produce the best quality. One evaluation criterion was Drywall Finisher 4's approval of a certain percentage. For evaluation, Drywall Finisher 4 walked across the partitions along with the team and meticulously checked for any imperfections, including raised, depressed, or over-sanded parts, by standing back and at times by zooming in or gently caressing the sanded seams and joints (Figure 14).[13] The ultimate feedback that distilled out of these inspections, later cross-checked with the engineering measurements, were the optimum sander setting and sanding paper grit, along with minor issues in finishing alignment and thickness, which would be fixed by revisiting the software, arm, and sprayer tip in the coming days.

To summarize, the salient points in robot–work integration at Canvas were Basil's reshaping of part of drywall finishing work into an iterative sequence of worker–machine action and the conventional workflow into work-sharing between finishers and the machine. On one side, these points show that the machine introduces a distinct way of working into drywall finishing. On the other side, they indicate the engineering team's grounded take on automation, by which they envisioned that the capabilities of tradespeople and the machine may combine to handle the complexity of the task. The other point is that the finishing quality appeared as a key element driving the development of the machine in the warehouse along with the acceptance and adoption of it by the industry. As I delineated in the chapter, the evaluation of quality is configured as an interdisciplinary practice, enacted by both the engineers with their measurement tools and numbers and the finishers with their trained eyes and hands. These different instruments are distributed among a pre-inspection to determine whether the drywall is ready for robotic targeted spray, a post-inspection to confirm that the sprayed parts are ready for sanding, and a final post-inspection to ensure that the finishing is up to industry standard. In the early-stage development of Basil I described in the chapter, it was the engineers' technical knowledge building on the finishers' explicit knowledge of the steps and tools and their tacit knowledge of quality that made the system possible. The 1K test once again showed the finishers' tacit knowledge of quality at play, this time keeping the system's performance in check and guiding its further improvement.

13 These imperfections in an actual job are manually fixed by adding or removing more mud.

Chapter 6

System Reconfiguring Construction Work: Canvas

6.1 Outline

By looking into the current technology design efforts at Canvas, in this chapter, I trace how Basil reconfigures drywall finishing work and discuss that it does so in two ways. First, drawing from an observation of a job site workflow in person and through the workflow visualization tool tracing worker-robot interactions, I show that the machine rearranges the set of skills needed in drywall finishing. This rearrangement consists of a degree of deskilling which makes a high-skilled part of drywall finishing more accessible for less-experienced finishers and a form of upskilling that brings in new intellectual and practical skills needed to work with the machine. Second, based on observations of the everyday work at Canvas and the network graph visualization tool mapping the work interactions around Basil, I show that the machine necessitates labor that involves distinct interdisciplinary interactions along with new roles and expertise attached to the drywall finishing trade. I discuss that the addition of technology designers as participants in the trade, the emergence of new roles in technology design for tradespeople, and the production of new knowledge through new interdisciplinary interactions are the changes that this labor at Canvas delineates for the trade. I conclude the chapter with a summary of the Canvas case.

6.2 An Ethnographic Study of Robotic System Development: Part 2

6.2.1 On and Off Site: Reconfiguring Drywall Finishing Work

Tracing Worker–Robot Interactions: Prototype B – Workflow Visualization Tool

Substituting the box-coating step with robotic targeted spraying, Basil incorporates a new and intriguing worker–machine feedback loop into drywall finishing, as exemplified by the 1K test. I considered that feedback loop an ethnographic site in which I could observe and analyze a new distinct software workspace to which part of drywall finishing was transferred. Within this workspace lies a demonstration of new skills that the machine brings into drywall finishing work. How a finisher operates in this workspace, what kind of interactions they have, and what new skills these interactions require were the questions I tried to explore to understand what is new and different about working with the system. In doing so, in addition to in-person observations, an important question became how to observe this workspace and trace the worker–robot interactions as part of finishing. The workflow visualization tool I developed provided a response to this.

The main data source for this prototype is the software that Canvas uses to record every second of any Basil UI screen in use and play these recordings for evaluation. With this, it enables an observer to see the UI during tests or jobs through the eyes of the worker and follow the worker's UI actions. The software also accommodates an around-the-clock data stream from various system components, such as the base, sensors, and CPU, and stores logs of these components. From this rich data pool, the UI recordings allowed me to conduct additional digital observations of interactions with Basil and the L4TS system during the 1K test and a job site pilot test. This data also inspired me to design a qualitative coding scheme, including the workers' UI actions and the machine's responses, to analyze these events. Lastly, recordings also allowed me to produce a large set of screenshots and, as a result of their analysis via the scheme, a text-based, detailed coding of the workflow (Figure 1). Prototype B, the workflow visualization tool that I developed, processes this text-based coded account and creates a map of the worker–machine feedback loop (Figure 1).

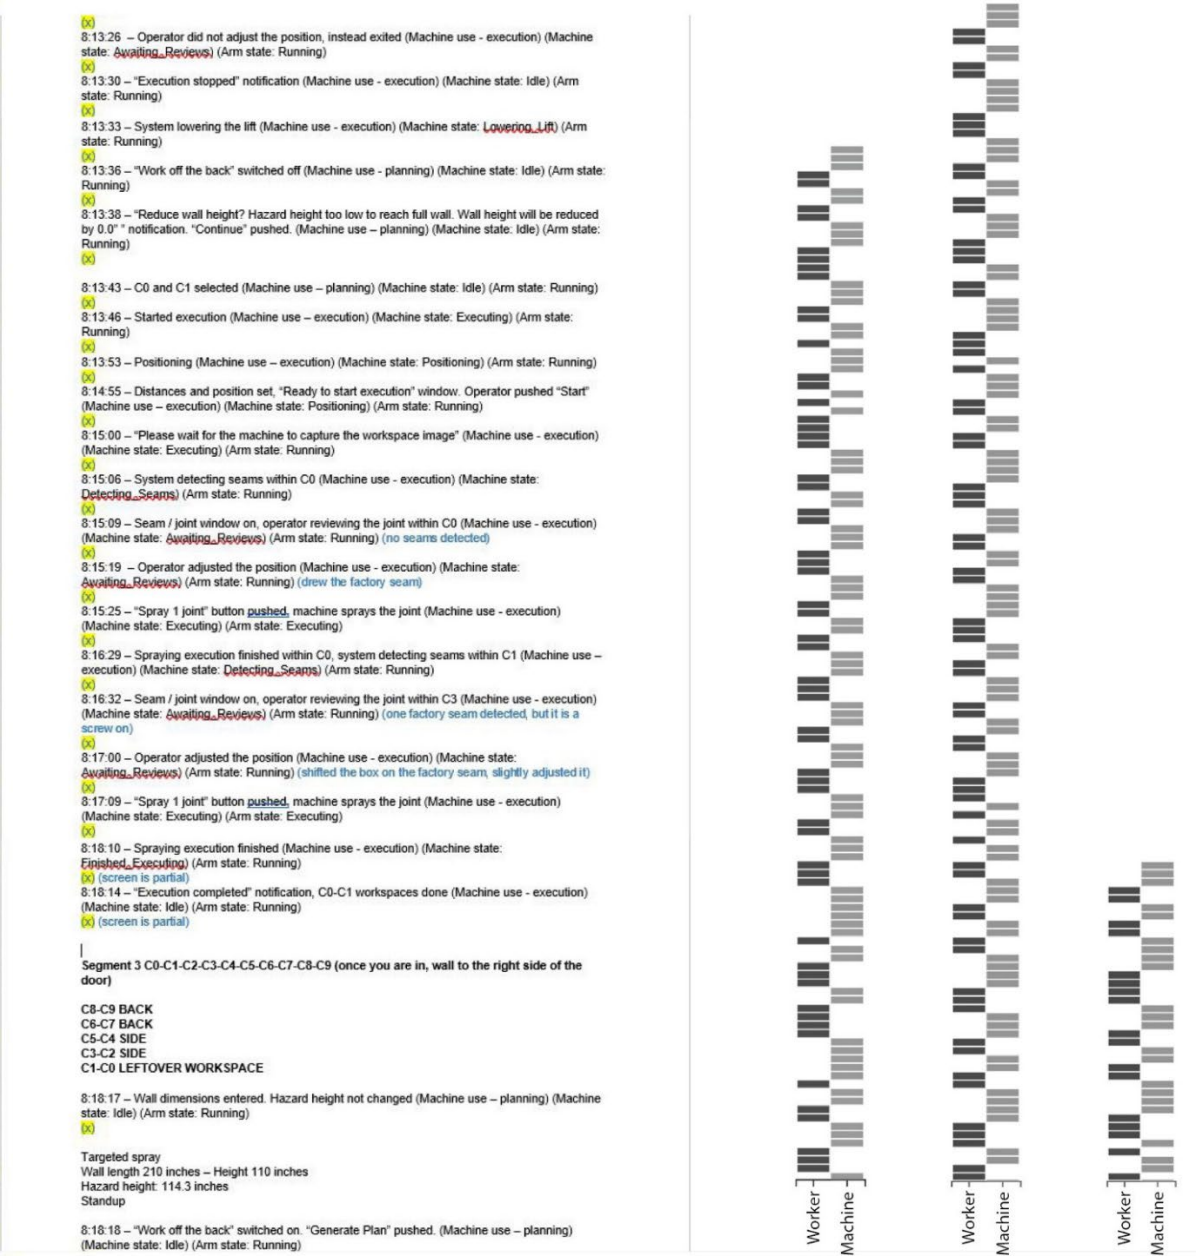


Figure 1. Sample coding (left) and a sample map representing a Canvas drywall finisher's UI actions and Basil's responses created by the tool (right)

The system was built using Python/NLTK (a Python library for natural language processing) as well as HTML, CSS, and JavaScript/D3.js (a JavaScript library for interactive data visualization).[1] Once the system is given the coding of the software workflow, it first scans through the words on each line and then labels each line in terms of the party in the interaction (worker or machine) and the type of action or response. It then stores these labels for each line in a dictionary and each dictionary in a list. As the next step, the system reformats this data set into a JSON file, which serves as input for

1 For more on "NLTK" by NLTK Team (2024), see www.nltk.org. For text analysis with Python, see Bird and coauthors (2009) and Bengfort and coauthors (2018). For more on "D3 (Data-driven Documents)" by Mike Bostock and Observable, Inc. (2024), see Bostock and coauthors (2011) or d3js.org.

the visualization component. The visualization component then generates the final map of actions and the system's output throughout the workflow. The position of each data point on the map refers to the actor and the order of action in the workflow. Color codes indicate the type of action or response. The map as a whole makes it possible to read the worker–machine feedback loop, which is initially converted into a nominal data set, in a spatial format.

User Interface Next to Pan and Knife: Rearranging the Drywall Finishing Skill Set

The L4TS system's first test in the wild was the finishing job of an office space at a prefabrication warehouse. Drawing from my in-person observations of this job site workflow, remote observations of it via the visualization tool, and interviews, I explored the skills that underlie the finishing workflow around Basil and L4TS, and that may become alternative or essential for finishers in the future.

Based on this exploration, I illustrate how Basil rearranges the set of skills that constitute drywall finishing work in this section. Rearranging the skill set is one way that I have found this newly integrated technology reconfigures work in drywall finishing. Upon closer inspection, this rearrangement involves a degree of deskilling in the sense that the machine provides the ability to less experienced or less skilled workers to handle a high-skilled step. However, the more visible aspect of this rearrangement is a form of upskilling, in which the machine updates the drywall finishing skill set by adding two main types of skills to the existing mix of conventional skills I described in the previous chapter. The first one is intellectual skills, which concern learning, adapting to, and job-planning with the machine towards a successful workflow. The second one is practical skills, which are needed to effectively work with and assist this technology. They are hands-on, tool based, and physical.

Having entered the Canvas nomenclature as the "job site pilot," this on-site test, which would happen over two days and include a wall area of two thousand to two thousand five hundred square feet, was designed to evaluate a list of criteria, including finishing quality, machine reliability, operator comfort with the UI sequence, and the level of the machine's adaptation to job site conditions. The goal was to maintain the performance achieved during the 1K test, which I described in the previous chapter. A relatively large team would be working on the job site pilot. Engineer 1, who coordinated the 1K test, would be coordinating the overall effort. Additionally, she would be documenting the issues that might arise, responding to them on the fly, and sampling thickness and fan width measurements on the site. Engineer 2, who handled the activity-time log during the 1K test, would be in charge of the photo and camera documentation, and Engineer 4, who observed the 1K, would be keeping the activity-time log this time. Engineers 5 and 6 would handle the pre-job inspection of the site and organize the transportation of equipment, respectively. The workflow would be the same as the 1K test, combining robotic targeted spraying and manual sanding with the vacuum sander. The difference was the wall configuration and operator. One of Canvas's finishers, Drywall Finisher 2, would run the machine this time. Along with the customer who requested the job, Drywall Finisher 4 would supervise, and once again, evaluate the final quality.

The job site pilot kicked off with Engineer 5 inspecting the site and confirming that the site conditions were suitable for deployment. He then passed the site location to the team for transportation and added an interior plan of the office space and wall specs into a shared coordination document. This plan drawing provided the team and me with a global view of the space, which consists of a workstation, a printer room, three identical office units, and a large conference room, all contiguous. Each room is a rectangular enclosure made up of three nine-foot-high, eleven-to-twenty-seven-foot

partitions, and a fourth one with an entryway, all of which may connotate qualities of space for architects and job scope for finishers. Based on this scope, the team allocated one day for transportation, on-site preparation, and wall prepping, two days for robotic targeted spraying, and another two days for manual sanding. Considering the quantity of walls to spray and sand, they also determined the inventory of equipment and consumables that would go along with Basil. Engineers 5 and 6, together with Drywall Finisher 2, packed the inventory, which included backup sprayer tips, coating and sanding tools, tower lights, and consumables such as mud drums and sanding paper. On the day after the inspection, the Canvas truck left the headquarters for the site with the machine and inventory in its cargo bed.

I arrived at the job site in the early morning on the day of the robotic targeted spraying. From the warehouse's industry-size parking lot, I entered the building and found myself underneath a grand timber frame that beautifully embraced a dull storage space, making it unexpectedly elegant. I entered the office area where the job would take place. The scene included several steel frames fitted with drywall panels hung in a stand-up layout, taped, and nail spotted, circumscribing the contiguous rooms on a large floor space (Figure 2). Three chairs with laptops, cameras, notebooks, extension cords, backpacks, and water bottles made up the pop-up observation station of Engineers

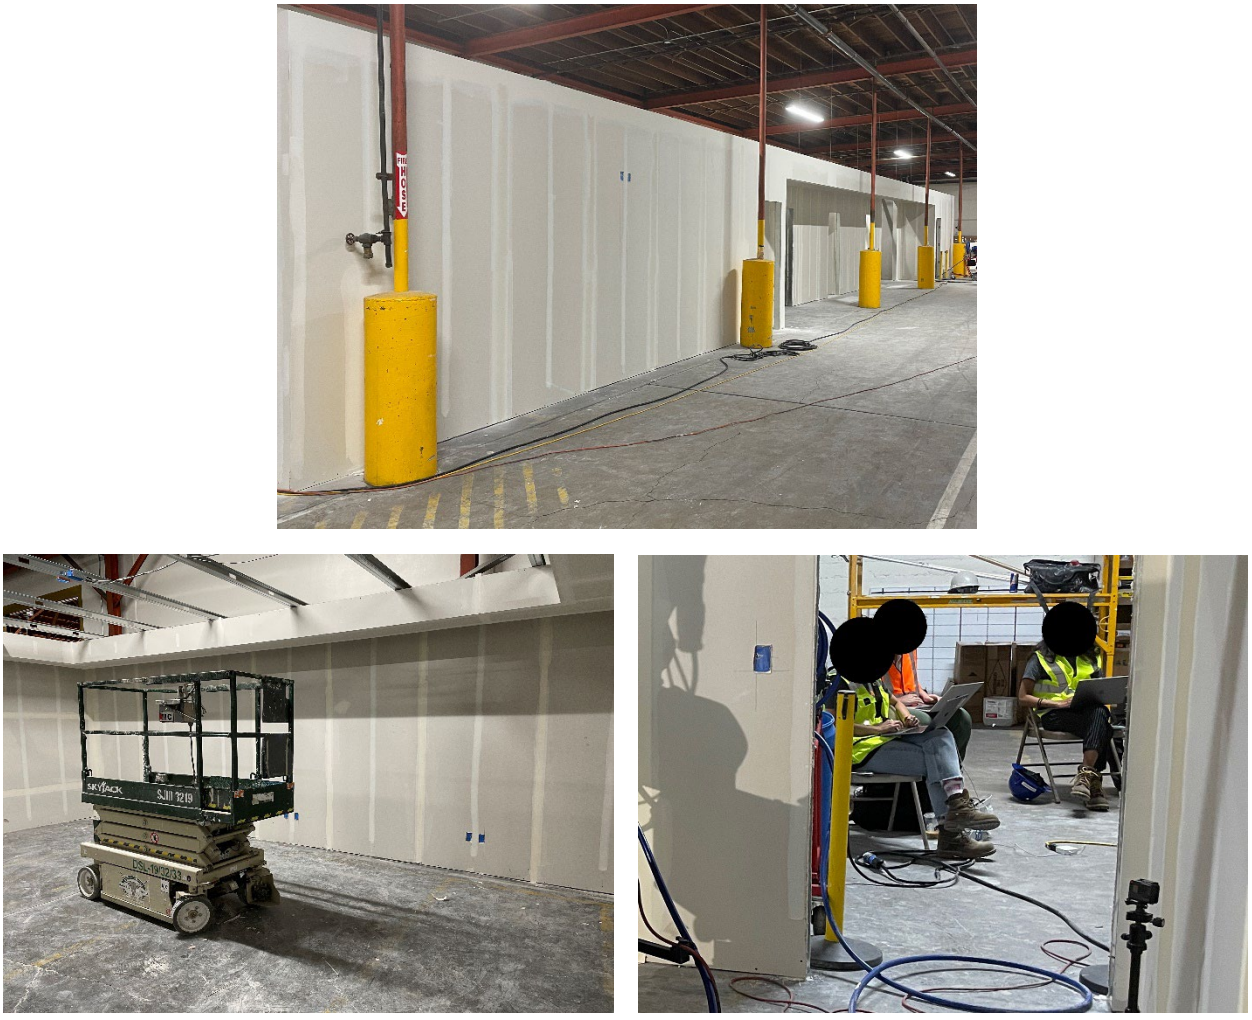


Figure 2. The office area (top), the future conference room (bottom left), and the pop-up observation station outside the future workstation (bottom right)

1, 2, and 4 just outside the future workstation, which would be the first room to be sprayed (Figure 2). I began my observation at this pop-up station. As Drywall Finisher 2 started with the first partition of the workstation, the active observation and recording of the engineers kicked off. The first partition was on the left side of the entrance and, therefore, visible to the engineers from the outside. Engineer 4 carefully followed Drywall Finisher 2's steps with the machine on the partition through this gap and documented each step with a short standard description on a designated spreadsheet. Engineer 2 monitored Drywall Finisher 2's actions on the real-time UI recordings on the software. She was also in touch with headquarters on a messaging app to provide regular updates. Engineer 1 was also on her screen at the station. She would soon go in to directly observe the workflow, check the outcome, and take interim measurements. As Drywall Finisher 2 was about to begin the third partition, I also went into the workstation to observe him and the machine closely.

The sequence that I observed in the room and later on the workflow visualization tool started with Drywall Finisher 2 entering the third partition's dimensions (Figure 3). He had just sprayed the last joint of the second partition from the side and could now begin working on the third one by spraying it from the back. Thus, he switched on the "work off the back" toggle and generated the wall plan. From the nine workspaces the system returned, he selected C8 and C9 within the first column and started the execution step. Using the controller, he adjusted the base's position aided by the reference lines on the UI. Once the distances were set, he ran the seam detection, and the system delivered an image of C8 on the review window. This time, the system could not mark the joint in the image, so Drywall Finisher 2 had to overlay the blue axis manually on the joint. He pushed the button, and the arm performed a single pass, covering more than half of the joint. Immediately after, the system lifted the base to C9, and delivered an image of the rest. Once again, the joint was not detected; at this point, I suspected that the relatively dark setting may be causing the issue. Drywall Finisher 2 added the axis, and the arm traveled along the joint and finished spraying.

Next were C6 and C7. Drywall Finisher 2 selected the workspaces and positioned the base. The sense-think-act cycle was then repeated. This part of the partition was better lit with a tower light he had the idea to bring in, so the system had no issues detecting the joint. For C4 and C5, he switched to spraying from the side and positioned the base accordingly. Once again, in two cycles with a single pass each, he finished the joint in these two workspaces. The third partition was completed after C2 and C3 because C0 and C1 were leftover workspaces featuring no joint. The same sequence was used on the fourth short partition with four workspaces, finalizing the workstation. The rest of the day covered the future printer room with nineteen workspaces on three partitions, one of the offices with twenty-seven workspaces on five partitions, and finally some blend-ins. Spraying also continued the next day with the two remaining offices and the conference room.

"Giving the artist [drywall finisher] a different tool" was Engineer 1's depiction of the workflow at the 1K and job site pilot tests I described. What is new when a drywall finishing job is planned around this new tool? What new skills might this new tool bring into the trade? I analyzed the job site workflow through the lens of these questions.

The major point that stemmed from this analysis is that the machine brings a degree of deskilling and a form of upskilling into drywall finishing work. Basil eliminates the box-coating steps of level 4 finishing as I noted in the previous chapter. Thus, the machine lowers the barriers to entry into this step for finishers with less experience or skills. This is a form of deskilling that happens inside the trade. More importantly, the machine introduces a different material process without box coating and way of execution with a software-hardware system, both demanding a form of upskilling to be able to work with it. Nonetheless, this upskilling does not render conventional finishing skills obsolete.

Bringing the drywall to the box-coating step, and being able to judge if the robotic execution of the step is successful, are still contingent upon conventional skills. New skills pertaining to the machine are added on top of these and update the skill set. Therefore, the baseline of this updated skill set comes from the trade, and the rest, which is new and needed to work the machine, must be developed on top.

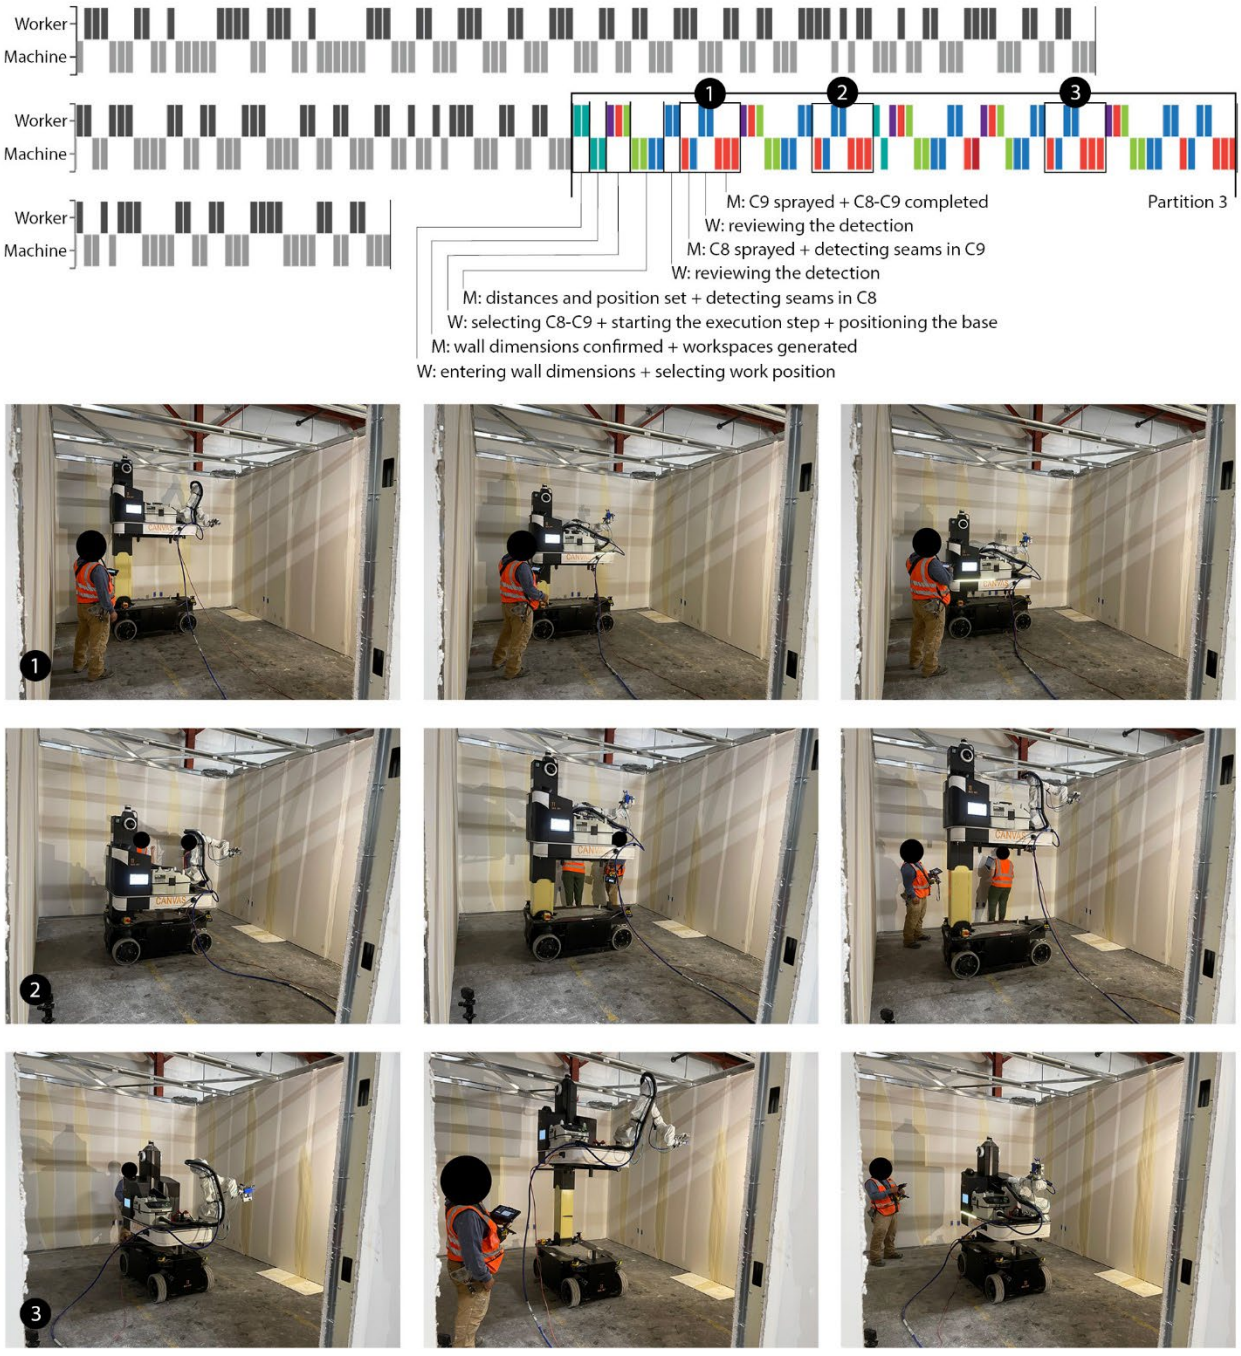


Figure 3. Visualization of Finisher 2's whole sequence with Basil in the first room (workstation) including the portion on the third partition color coded (top), and moments of spraying the joint in C8-C9 from the back (top row), the adjacent joint in C6-C7 from the back (middle row), and the adjacent joint in C4-C5 from the side (bottom row)

A couple of key characteristics of drywall finishing explain why a combination of trade and machine skills is needed to successfully accomplish a machine workflow. First, finishing is distinct from other more mechanical drywall tasks, such as drywall installation or framing, in that it is a material-oriented coating task. Conventional drywall finishing requires rounds of coating and quality evaluation to take the partitions to the specified levels as I explained in the previous chapter, whereas hanging drywalls may not necessarily require as much technical standard or aesthetic quality evaluation. Thus, visual and tactile quality inspection are inseparable from drywall finishing, even when it is delivered with the machine, as in the 1K test. Second, certain steps of finishing demand "a fine eye for detail or more dexterity," as General Manager 1 put it, and handling these steps is beyond the machine's current capabilities. For instance, seams and joints still need to be taped with the bazooka before the robotic targeted spray. Tighter enclosures, corners, soffits, and areas around pipes that the machine cannot access or would spray too coarsely must be done manually with the pan and knife. Thus, to complete the targeted spraying step, an individual finisher or a team still necessitates at least apprentice-level skills in mud preparation, taping, and nail spotting as well as basic knowledge of hand tools. To be able to complete the entire job with the machine, this should be accompanied by foreman-level skills for pre- and post-inspection of quality along with subsequent manual fixes.

The job site pilot team possessed all the experience levels needed to bring in trade skills next to machine skills.[2] On the day of the machine's transportation to the site, Drywall Finisher 2, a journeyman finisher, performed the pre-inspection of drywalls, which had been taped and nail spotted by the customer's own drywall finishing team. He spotted some minor faults on the drywall panels as well as a few dents on the taped joints, fixed them by adding more mud with his pan and knife, and finalized the initial wall preparation for the machine.[3] On the days of the robotic targeted spraying, Drywall Finisher 2 was in charge of preparing the mud. Having done this numerous times, he knows by heart the water/joint compound ratio and how long to stir the mix to reach the right viscosity. He also operated the machine and performed blend-ins with his pan and knife on spots that required it on those two days. On the next day, Drywall Finisher 4 evaluated the mud thickness, the consistency of thickness, and the texture on the sprayed joints using her visual and tactile inspection skills and made sure that they were ready for sanding. Drywall Finisher 2 then fixed the imperfections they found via manual touch-ups with his pan and knife. On the same day and the one after, he completed the sanding with the vacuum sander. As the final step, Drywall Finisher 4 came back for the final post-inspection, and Drywall Finisher 2 once again fixed the imperfections by adding or sanding more mud.

What are the new machine skills that combine with these? Observing the job site pilot team, interviewing its members and the other engineering team members who prepared the L4TS system, and examining Drywall Finisher 2's job site workflow through the visualization tool enabled me to document the new intellectual and practical skills pertaining to the machine.

The foremost intellectual skill that I traced is being open to using a new tool and agile in learning new skills. The finishers who have the chance to use Basil in its early days are the "early adopter[s] [of the system]," as Engineer 12 put it in our interview. "There's a lot of things that they [finishers] have to keep up with, in order to stay on top of what our system is," she said to emphasize learning as an inherent aspect of working with the machine. "Patience and persistence" in this learning process are important, as General Manager 1 affirmed. She mentioned that those traits vary across finishers at different levels of experience. Training Manager 1, who trains finishers to operate the

2 Alternatively, an individual finisher may also have both the trade and machine skills in varying levels. For instance, both Finishers 1 and 2 are able to operate the machine, fix imperfections, sand, and to a degree handle pre-inspection before these steps and post-inspection after.

3 Such fixes continued to take place here and there as more faults were encountered by Finisher 2.

machine, confirmed these differences and highlighted a few specifics related to demographics: he explained that more experienced finishers stand “in the middle” of excitement and skepticism towards the machine, while finishers who are new to the trade “tend to find the machine exciting.” More direct feedback about the importance of being open-minded to learning came from Drywall Finisher 2, who operated the machine at the job site pilot. He recalled how he navigated through different jobs and skill sets in his career, including baking, sewing, and carpentry, until he learned drywall finishing. His open mindset was the starting point of his mastering the use of Basil.

Another intellectual skill is knowing how to adjust one’s knowledge and ways of performing drywall finishing to adapt to the new steps Basil brings into the task. As I pointed out in the previous chapter, the main difference between robotic and conventional drywall finishing is the box-coating step, which is eliminated by the use of Basil. Instead of applying three layers of box coat on the tape with drying time in between, a finisher applies one thick coat with the machine without waiting for the tape to dry. This “wet-on-wet” spraying process and the amount of mud applied to the partition are “new to workers, and different than what they know,” said Training Manager 1, as he shared his observations during training sessions with finishers. This “learn[ing] a new way of approaching drywall,” as Engineer 12 called it, is critical for knowing the differences and variability in outcomes between targeted spraying and using boxes. Thus, it is critical for managing quality control. During the job site pilot, it was Drywall Finishers 2 and 4 who had the eye and knowledge necessary for reading the quality of this new type of outcome. Both finishers knew when to “touch it up, add more mud, sand more” and when to “walk away,” as Drywall Finisher 4 put it.

The machine’s integration also turns part of drywall finishing from mobile work with hand tools into more stationary and observational work with the UI, which is clearly an adjustment for finishers. An indicator of this is the repeating sequences of inputs Drywall Finisher 2 had to provide to the system, which are seen in the visualization that represents the entire job site workflow in the first room (Figure 4). The colored portion highlights his various recurring inputs while spraying the third partition, including positioning the machine in front of each column, attending to the feedback loop during the seam detection step of each sense-think-act cycle, and reviewing and adjusting the output if necessary (Figure 4). The same visualization also shows the machine’s responses or spraying actions in return, which Drywall Finisher 2 still had to pay attention to or monitor in most instances (Figure 4). The feedback that Training Manager 1 has received from his trainees puts into words what the visualization shows: this repeated interaction with the UI makes finishers, who are very much used to working by moving, start to think that their job becomes to “tend to the machine.” However, I observed finishers with more experience with Basil finding ways around this. Doing concurrent manual work during spraying periods, such as short wall preps or cable organization like Drywall Finisher 2 did on the job site, is a good example of how a finisher can skillfully adjust their flow according to the machine.

The last intellectual skill that I traced is knowing how to plan the drywall finishing job when the machine is part of it. More specifically, this concerns planning the steps and the space in which they unfold at the job and wall levels. At the job level, this falls under the foreman’s responsibilities. Careful scheduling of the steps to maintain finishing productivity with the machine and safety by not getting in the way of other trades is important for these professionals to handle. The first key skill for productivity is planning the sequence appropriately, in particular, knowing how to “mesh” the conventional steps with the machine step, as both Engineer 7 and General Manager 1 described. Construction inherently involves contingencies and improvisations that might affect the sequence, but the ideal strategy set by Drywall Finisher 4 for Canvas is to run and complete the machine step at once. As the job site pilot demonstrated, this strategy involves finishing the taping and nail



Figure 4. Instances of Drywall Finisher 2 providing inputs for positioning the base (p) (green), reviewing seam detection (r) (blue), and monitoring spraying (m) (orange)

spotting of all the partitions on a floor first, as opposed to concurrently operating the machine as the taping and nail spotting progresses.[4] The other key skill for productivity is "keep[ing] walls more or less [on] the same schedule," as General Manager 1 pointed out. This means dividing the labor so that work progresses at the same pace on the wall parts that the machine cannot handle and on those processed by the machine. This was not a concern at the job site pilot because the partitions were simple and not very detailed.

At the wall level, planning the sequence is the responsibility of the finisher running the machine. Similar to scheduling the drywall finishing steps and the space for a foreman, the ability to effectively plan the actions with the machine and the space is important for finishers to maintain productivity and safety. The key is to minimize movement. Part of this is done by arranging the work zone, as I saw Drywall Finisher 2 do during the job site pilot. For instance, before starting the workflow, he put the extension cord for the machine in an appropriate spot to avoid repositioning it every time he moved to a new room. When working in the small rooms, he placed the mud drum outside to allow space for conveniently maneuvering the machine and minimize the movement of this equipment (Figure 5). At each room's entrance, he placed neon-yellow barricade tape to keep people out while the machine was in action (Figure 5). The other part of minimizing movement is about arranging the sequence of worker–machine actions, which Drywall Finisher 2 also demonstrated nicely during the job site pilot test. A practical approach he followed was to measure all the partitions in a room as he entered it rather than after each machine positioning. He penciled the numbers on the partitions and entered them into the system once he reached a new partition. To minimize the machine's positioning, he reviewed the layout, angles, and lengths as he stepped into a room and accordingly

4 In this way, partitions remain taped without any coat for some time into the job, and not seeing concurrent progress becomes stressful for foremen at times, as General Manager 1 pointed out. Therefore, "trust[ing] the machine will deliver" as she emphasized, is crucial.

devised the optimum spraying order from the machine's side and back. Going clockwise, he sprayed the last joint of each partition from the side and, without making a turn, the first joint of each following partition from the back (Figure 6).

On the other side, the practical skills needed to work with Basil are the ones finishers are trained in and practice on sites, as I categorized earlier in this section. Some of these are the very basics of the robotic targeted spray workflow, such as how to install and clean the sprayer, drive and position the base, spray around wall penetrations, do blend-ins between the hand-coated and targeted-sprayed areas, and sand with the vacuum sander. However, the ones that I would like to reflect on are the skills brought about by the work-sharing between finishers and the machine, or, more precisely, the worker–machine feedback loop I described earlier and visualized in this section.



Figure 5. Finisher 2 laying out the extension cord, mud drum, and machine before starting a room (left) and the neon-yellow band to keep clear (right)

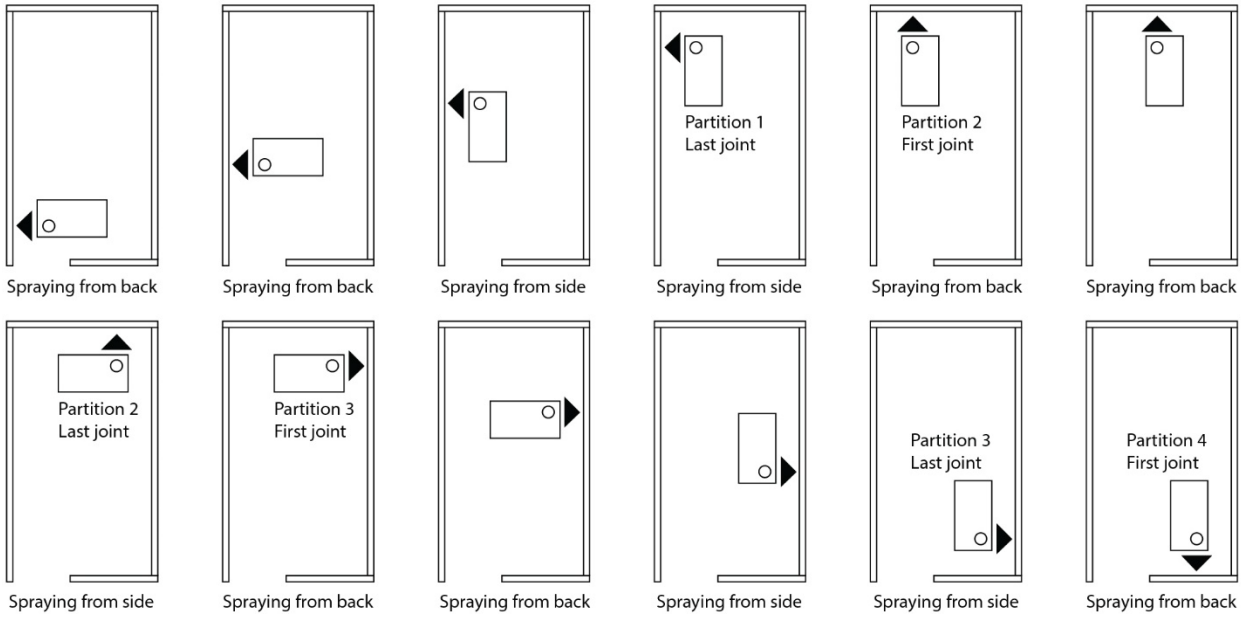


Figure 6. Finisher 2's sequence inside the workstation—spraying the last joint of each partition from the side and the first joint of each next partition from the back

The foremost practical skill I traced on the job site and through visualizations is operating in both software and physical space. Adapting design scholar Malcolm McCullough's (1996) hand-eye-tool description to drywall finishing, we may say that conventional finishing involves the "hand, eye, and tool converg[ing] in one place," namely, the partition.[5] It is "the hand [with a tool] work[s] a material," which is mud on the partition, while "the eye follow[s] it continuously" (McCullough, 1996). Robotic drywall finishing brings about, in McCullough's (1996) words, a "hand-eye separation," but it occurs differently in his example of using CAD software, in which "the hand moves a mouse" providing digital input, "while the eyes look at a screen" following the actions and changes in the outcome in software space. With Basil, the hand performs the control and multi-touch providing the sequence of inputs that the visualization shows, but the eye follows the machine's responses, spraying actions, and changes on the wall simultaneously in software and physical space (Figure 7). This new hand–eye configuration was at play when Drywall Finisher 2 would enter the partition dimensions, position the base, and review the seam detection, which I observed on the job site and then on the workflow's visualization and photos (Figure 7). In each of these, the system's UI provides a digital spatial view that was not available to the finisher before (Figure 7).

For instance, every time Drywall Finisher 2 entered the partition dimensions and generated the plan, he translated the measurements he took in physical space into the column workspace grid, which is a spatial data representation in software space. Embedded in this representation that he was repeatedly checking is information about what will happen in physical space, such as columns and workspaces showing how many times he should reposition the robot or how much of the partition he completed. As an experienced operator, Drywall Finisher 2 is fluent in this data-enabled way of working, which entails dealing with information and a partition simultaneously (Figure 7). He knows how to read the partition computationally through information. Similarly, when Drywall Finisher 2 was positioning the base, he was using the controller with his hands while monitoring the results of his input, including where the robot stood or how far it was from the partition, in both the UI's Cartesian space and physical space (Figure 7). When he reviewed the seam detection output on the UI, he sometimes adjusted the alignment of the detection, which determines the arm's spraying waypoints, with his fingers and then tracked the arm's movement and change on the physical partition (Figure 7). Such inputs into the system can also be impromptu, as Drywall Finisher 2 exemplified on the job site. He called this "tricks and ideas you learn on site" and cited using the tower light to illuminate the scene as an example. As he explained, he modified the physical environment to "help the machine" generate better images and execute the detection.

Lastly, evaluating the machine's outcome is perhaps the most important practical skill. As General Manager 1 stated, "the machine can't know that there's a sub-millimeter-level part that's out of plane," and this is why quality evaluation skills are still needed, as I discussed earlier in this section. On the job site, Drywall Finisher 2 performed the pre-inspection of drywalls, visually scanning and at times gently running his fingers on tapes to see or feel the depressed parts. He reviewed nail spots to verify that they were well-coated. After the targeted spray, Drywall Finishers 2 and 4 carried out the post-inspection and checked whether there was enough mud, consistent thickness, and no high or low parts on joints, whether sprayed and hand-coated areas were well blended, and whether the desired stipple spray texture had been achieved. The two conducted the same inspection after sanding by looking at and touching the parts they suspected of being uneven due to under- or over-sanding. There is no rule or universal criteria in this regard. They know how to identify imperfections based on an intuitive sense of quality resulting from their trade experience. Most of the time, they also know why a particular imperfection occurred based on their experience with the machine, which

5 See McCullough's (1996) chapter titled "Human Context" in his book *Abstracting Craft: The Practiced Digital Hand*.

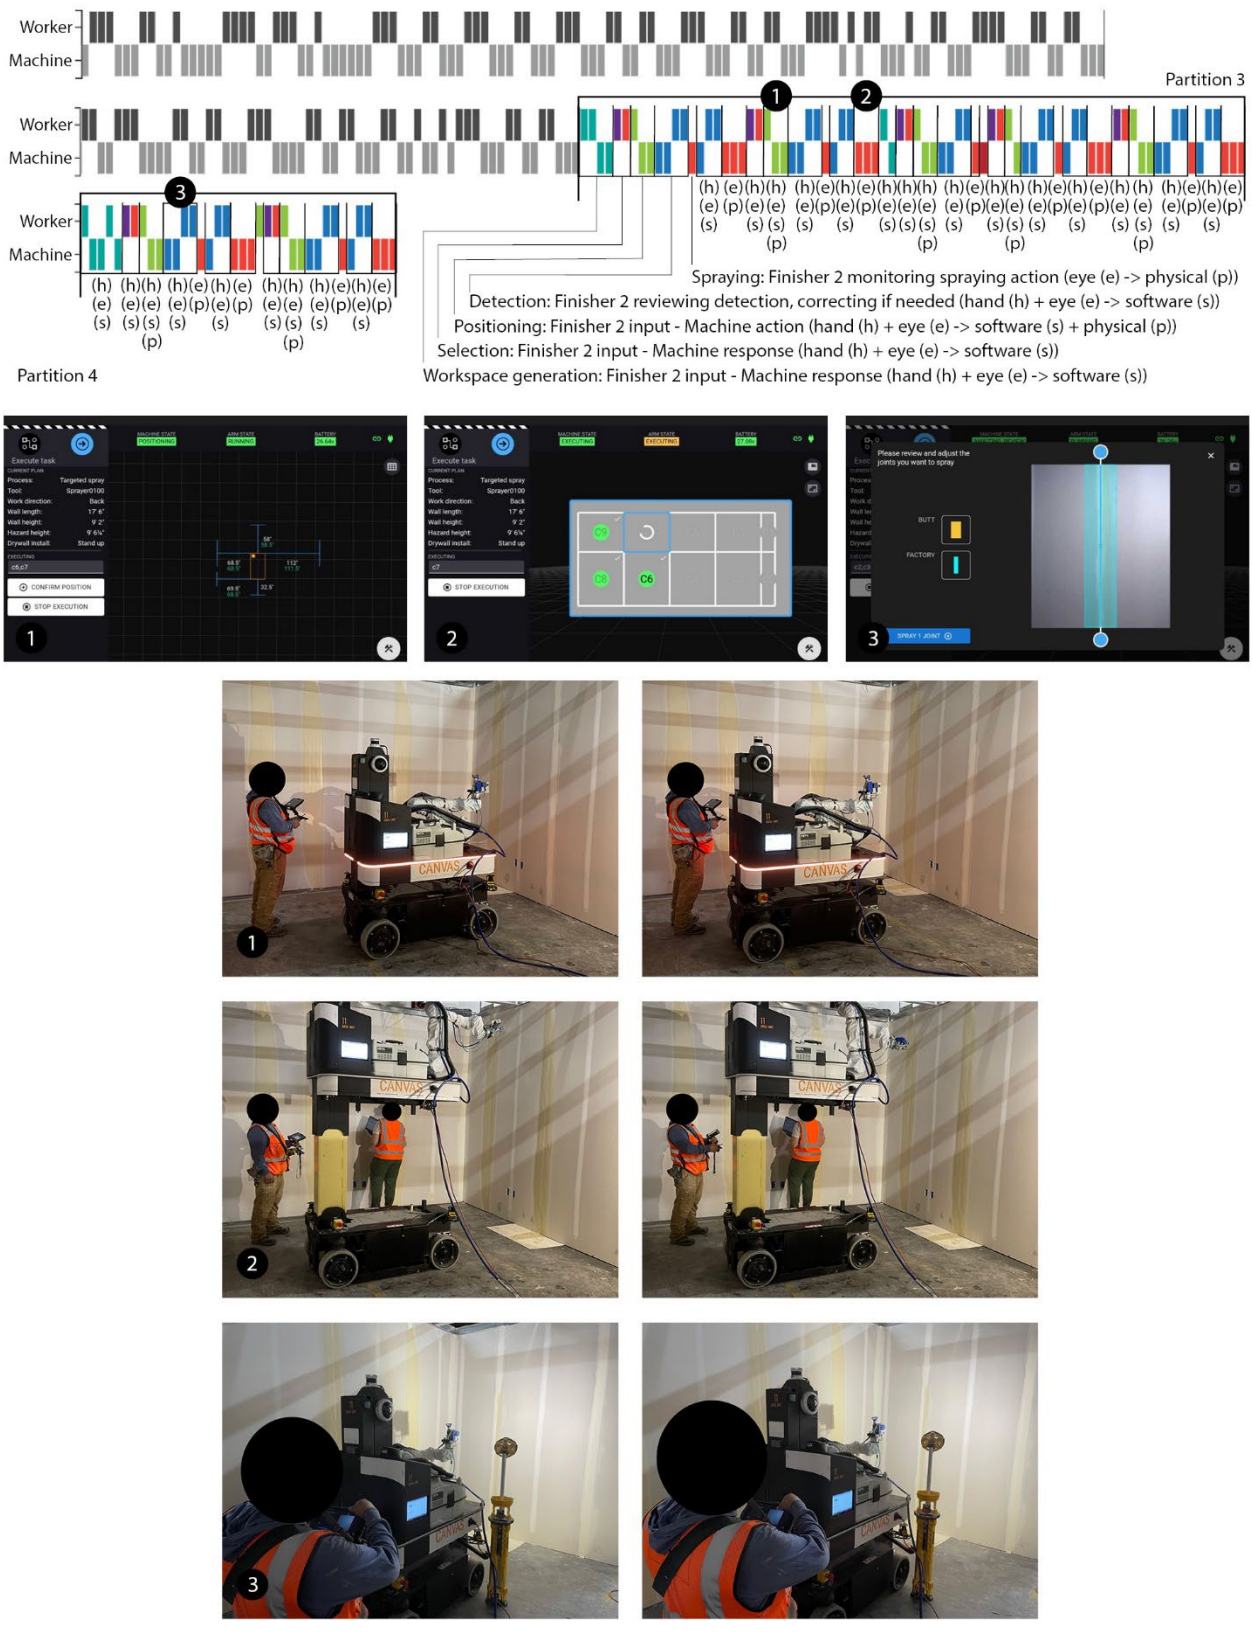


Figure 7. Instances of Drywall Finisher 2 operating in software and physical space while positioning the base (hand and eye) (green), reviewing seam detection (hand and eye) (blue), and monitoring spraying (eye) (orange)

is new. When performing box coating, finishers get “physical feedback”—as Engineer 1 put it—from the wall in real time as they alter the partition. It is the finishers’ vision and dexterity that enables them to have control over quality. When working with the machine, they do not have this feedback or direct control, so they need the skill to understand the reason behind a particular spraying outcome and how to fix it with manual touch-ups, as Drywall Finishers 2 and 4 demonstrated on the job site.

In sum, Basil updates the drywall finishing skill set by adding new skills to the mix. Drywall finishing is a material-oriented practice with standards for outcome quality. Therefore, performing it with the machine still necessitates the conventional finishing skills based on the knowledge of the trade and the dexterous way of practicing it manually, combined with new skills specific to working with the machine. This mixed skill set is sometimes distributed among a team of finishers, as I showed in the job site pilot, or within one worker, as Drywall Finisher 2 epitomized. In any case, the targeted spray workflow still involves the pan and knife alongside the machine’s UI. As I outlined in this section, the new skills that Basil brings to drywall finishing are intellectual and practical. The targeted spray workflow requires understanding the differences in the way the machine produces the outcome. Therefore, it adds agility in learning a new tool, adjusting to the idiosyncrasies of drywall finishing with the machine, and planning the finishing steps at the job level and the actions with the machine at the wall level as new intellectual skills to the conventional skill set. The workflow also demands, to borrow from McCullough (1996), a “practiced digital hand” from finishers to operate in the feedback loop with the machine. It introduces operating in both the software and the physical world to assist the machine, as well as evaluating the machine’s outcome, as the new practical skills to develop.

Mapping Robot Development Work: Prototype C – Network Graph Visualization Tool

The Canvas study focused on two types of work as objects of analysis: construction work, more specifically drywall finishing work at the core of the robotics practice, and the system development work the robotics practice conducts on a daily basis. Analyzing the latter was important for understanding the collaborative work that is carried out around the system and, perhaps, becoming part of the construction industry. How to make this unfamiliar organization familiar, comprehend and analyze its complexity, and make this analysis communicable was a challenge and a reflexive question I asked myself as a researcher. The engineering team I engaged with offered a large group of participants who consented to participate in the study and an intricate network of work interactions that I identified as key to the development, test, and deployment operations. I began to trace these through in-person observations and interviews. To support this effort with a remote high-level analysis and more granular documentation, I developed a prototype network graph visualization tool mapping these interactions.

The data that became the input for this prototype was created by reformatting the meeting information on the Google Calendar app, which is the backbone of the engineering team’s work planning and, therefore, documents their interactions. Reformatting involved manually writing and organizing the calendar data, represented in a typical calendar grid, into a data structure. Ultimately, ten weeks of meeting information were reformatted into ten sets of nested dictionaries in Python, each corresponding to a week. Each parent dictionary in the set is dedicated to one participant (Figure 8).[6] Each child dictionary represents a meeting or an activity that the participant attended (Figure 8). Meeting or activity names usually contain the name of a system component or an operation, so they convey themes or what the meeting interaction was about.[7] Each of these child dictionaries contains a list of the other participants in the meeting and who engaged with the

6 Participants are anonymized with pseudonyms.
7 Meetings described in the chapter are anonymized with pseudonyms.

discussed theme—a particular system component or operation (Figure 8). Prototype C, the network graph visualization tool that I built, is a combined data processing and visualization system that generates a network graph visualization of engagements among participants and between participants and themes.

The system was developed by using Python/Pandas (a Python library for data analysis) as well as HTML, CSS, and JavaScript/D3.js.[8] It handles the data processing in two layers. The first layer takes the raw data (a set of nested dictionaries) as input, computes the number of each participant's engagements with every other participant and every theme in a week, and generates an adjacency matrix that contains the participants, themes, and number of engagements. The second layer takes the adjacency matrix as input, sorts the participants from the themes and by teams, and reformats it into a JSON file that contains the nodes and links data of the graph. Based on this data, the visualization component generates an interactive graph for each week, which displays the participants as colored nodes and themes as black nodes, and the participants' number of

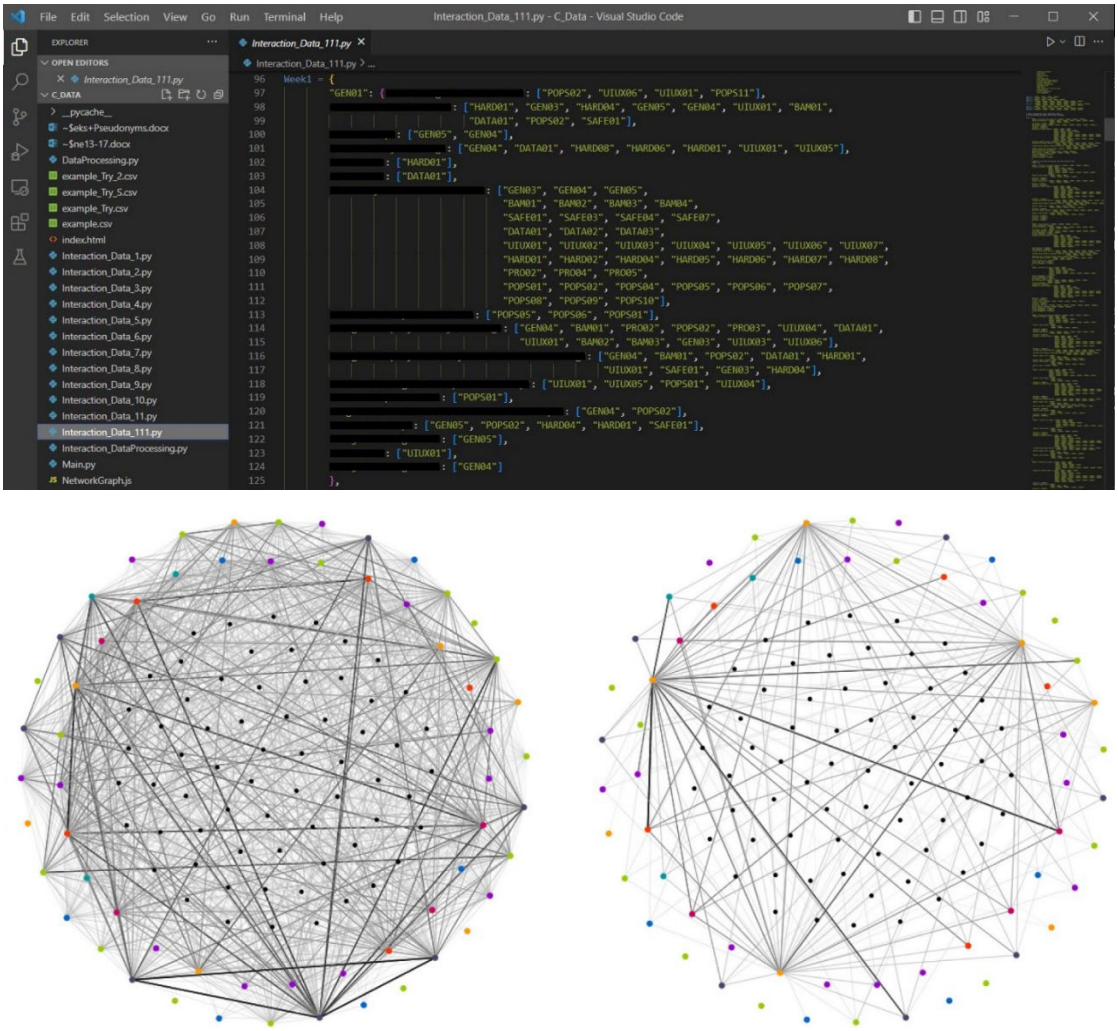


Figure 8. A nested dictionary (top) and graphs showing all participant interactions (bottom left) and a team's interactions (bottom right) in a week

8 For more on "Pandas" by the Pandas Development Team (2024), see McKinney (2010) or pandas.pydata.org. For data analysis with Python, see McKinney (2013).

engagements with each other and with themes as links of variable thickness. Each participant node is color-coded according to their team, and a mouseover on that node shows the participant's pseudonym. The network graph presents all interactions as well as the interactions of a specific team or participant in a given week (Figure 8).

Labor at Canvas: New Participants, New Roles, New Knowledge

Looking into the daily work of teams behind a robotic system reveals aspects of technology design labor that shapes it. Through the network graph visualization tool offering me a map of Canvas's engineering team organization, observations of meetings, and conversations with members, I took a distant and close look at the labor that goes into Basil's development and deployment.

Based on an analysis of weekly interactions I mapped via the tool and the meetings I engaged, I discuss the two important elements that I traced within the labor at Canvas in this section. The first one is the interdisciplinarity and diversity of expertise including various engineers and workers informing the development of Basil. The second one is new forms of labor, in other words, new roles and expertise that are essential to operate the robot and becoming intertwined with the drywall finishing trade. Thus, next to rearranging the skill set, a second way Basil as a new technology reconfigures work is by introducing this body of distinct interdisciplinary interactions and new expertise into the trade. Specifically, the change this new technical labor at Canvas delineates for drywall finishing comprises three aspects: the addition of technology designers as participants contributing to the trade, emergence of new roles in technology design for tradespeople, and production of new knowledge for the trade through distinct interdisciplinary interactions.

As the larger of two main teams—the other one being the sales and marketing team—the engineering team is structured in sub-teams focusing on software, hardware, the robotic drywall finishing process, training-test-deployment, and management.[9] The Software Team, which I will call Team S, consists of four specialized teams with members coming from mechanical, robotics, and software engineering. Each is responsible for the development, testing, debugging, and documentation tasks related to mobility and manipulation (B), sensing (A), user interface (U), and data stream and simulation (D), respectively. The Hardware Team, or Team H, is composed of mechanical and electrical engineers, electrical technicians, and fabricators. These members carry out hardware design, fabrication, purchasing, and testing. The Process Team, or Team P, comprises members from the sciences and mechanical engineering. This team's role is to understand and model the relationship between software-hardware, mud, and finishing quality through experimentation and testing, which has been central to designing the robotic workflow in the early years and, more recently, to its further improvement. The Operations Team, or Team O, includes mechanical and civil engineers, technicians, managers, and the in-house drywall finishers, who share responsibility for machine performance testing, operator training, job site deployments, support, and maintenance. Lastly, Team G is not an official team but a set of general managers, specifically mechanical and robotics engineers with project management experience and former construction professionals who oversee and make decisions with the sub-team leaders.

Based on initial visualizations via the tool and observations, I quickly noticed that the work is not confined to the sub-teams and is often conducted by a mix of members gathering around a theme in meetings. Every week saw more than sixty of them, generating a dense network of interactions within and across teams about coordination, system development goals, system issues, tests,

9 I did not focus on the sales-marketing team in this study.

deployment, and project management (Figure 9-A). Seeing all these meetings or themes and the participants that convened and collaborated around each theme in one week in one compact visualization provided me with a high-level view of the organization. After a few weeks of examining this in parallel with conducting observations, certain meetings or themes began to stand out to me, demanding more attention. Filtering the interactions by teams revealed the centrality of some participants, who, if this was verified by the observations, would become the persons with whom to engage further (Figure 9-B).

For instance, a theme that appeared notable was the 1K test I described in the previous chapter. I learned about it from Engineer 1, who seemed to be a central figure. Noticing the number of weekly meetings related to this test, which eventually reached more than seventy in a three-week span, I decided to follow it closely. Isolating each meeting in the visualizations indicated that the conversations often featured an interdisciplinary mix of members (Figure 9-C). A salient demonstration of that was Meeting F, an early morning meeting about finishing quality tests. Coming from teams P, U, and O, the participants in this video conference meeting included Engineer 7, a process engineer with a background in the sciences and engineering, Engineer 1, an operations engineer with a background in mechanical engineering who would later orchestrate the 1K test, Engineer 8, a user interface/user experience (UI/UX) designer with a background in design and mechanical engineering, and Drywall Finisher 4, who would later conduct the final quality control after the 1K test (Figure 10).

Engineer 7, who was leading the meeting, went calmly over a list of items about the L4TS system, which would soon enter its final round of tests, including the 1K. Each item was a prompt to hear from the other participants. The first item concerned the system's validation by the crew. Engineer 1 reported, in a casual and confident tone, that the finishers were still getting used to the system and evaluating it to give feedback. Engineer 8 followed and proposed to obtain the finishers' feedback during the interviews that he would conduct with them on other UI-related matters, such as the clarity of the naming convention used on the UI widgets. Meanwhile, Drywall Finisher 4 was not fully satisfied with the crew's delivery quality with the system so far. She recommended more training. Engineer 1 assured her that there would be additional training on L4TS best practices when the software release would be fully ready after minor changes. The second and the last items were related to the preparation of the testing area for the 1K test. The team was happy to hear from Drywall Finisher 4 that four partitions amounting to the necessary square footage were being drywalled.

Another theme that stood out from the visualizations and observations was the job site pilot I covered earlier in this chapter, which was the product of several meetings over seven weeks. For Teams S and P, resolving the misalignment of the spray that appeared in the 1K test was the main preoccupation in those weeks and a primary theme of several meetings that gathered diverse technical expertise towards solving this issue. One of these was Meeting S, an internal Team S meeting including members from Teams B, A, U, D, and G (Figure 11). With part of the members attending online and others from a video conference room in the warehouse, this meeting was marked by a mix of concern and confidence.

After a list of updates on the robot simulator building efforts, development work schedules, and the new release, General Manager 5 provided updates on the misalignment fix. With a clear and confident tone, the robotics engineer, who has extensive expertise in robotics, explained the two-part solution that the team had been working on. As General Manager 5 reported, describing the first part

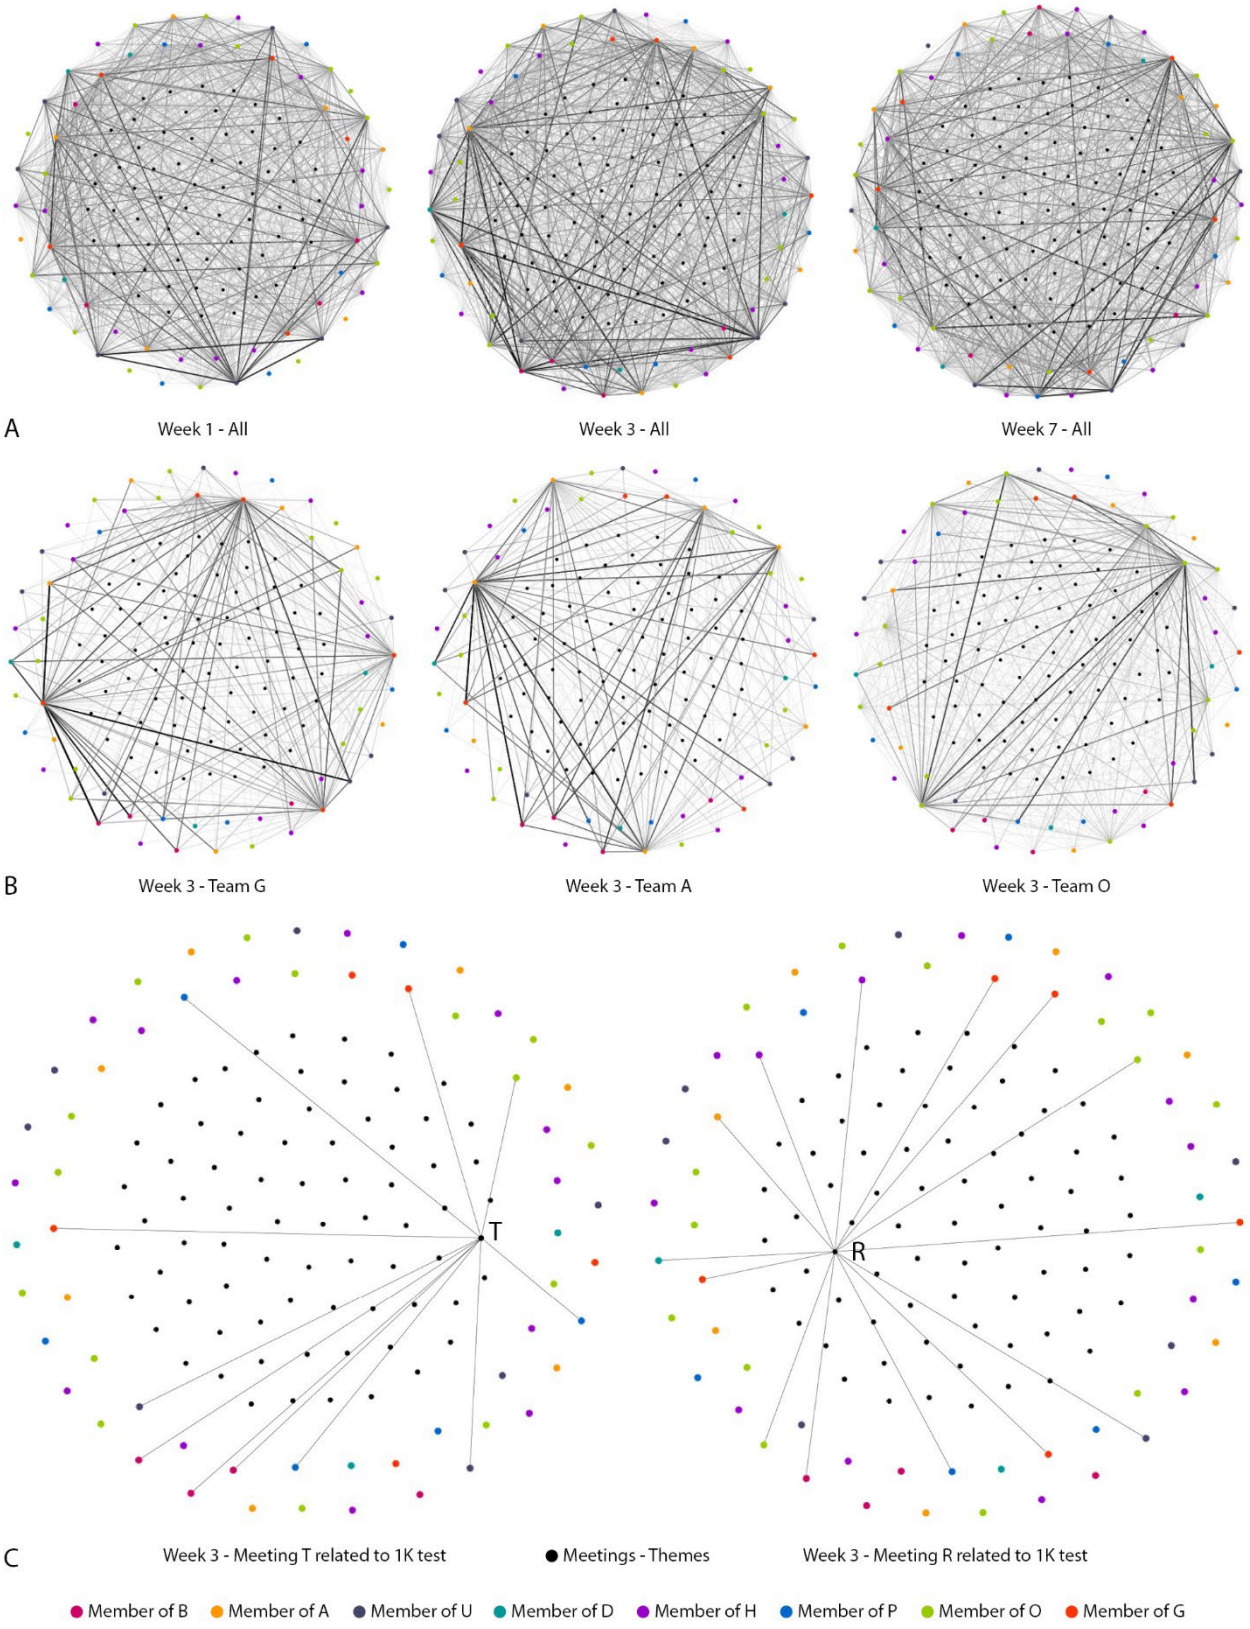


Figure 9. Weekly interactions of the whole engineering team (row A), interactions of teams G, A, and O in week 3 (row B), and two of the 1K test meetings attended by an interdisciplinary mix from different teams (row C)

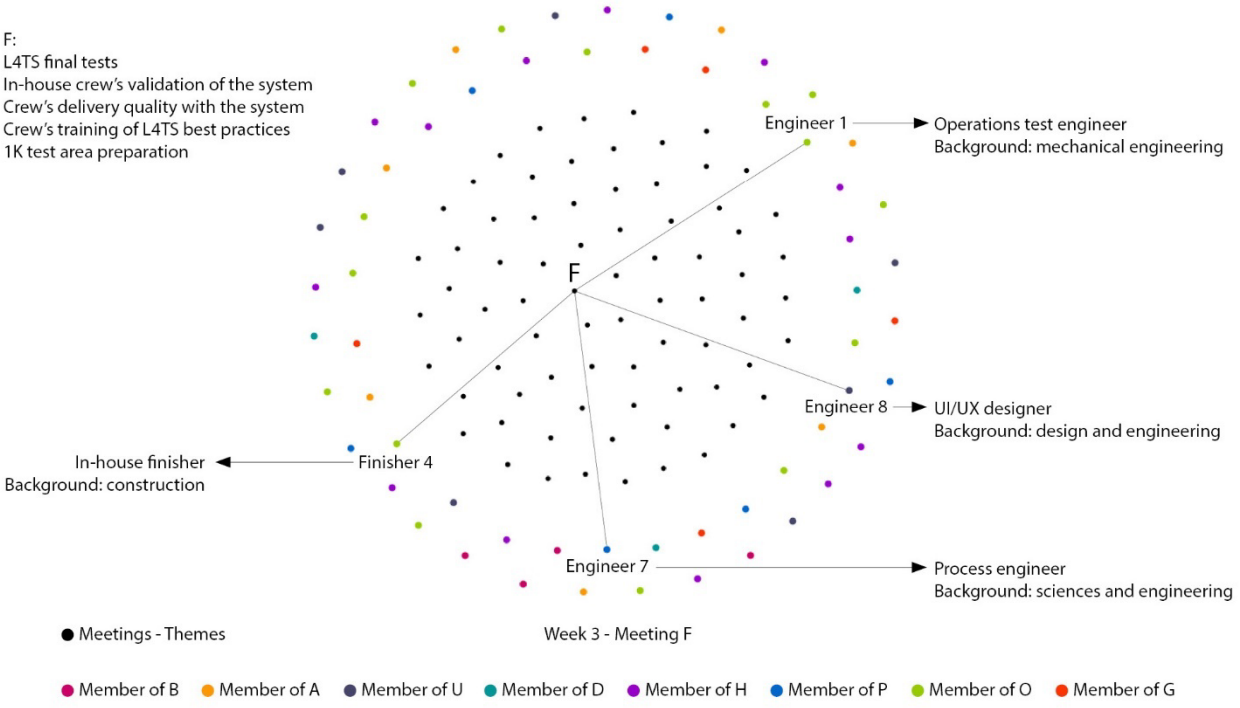


Figure 10. Meeting F showing the mix of backgrounds informing the system: in this case, validation, delivery quality, training, and the 1K test related to L4TS

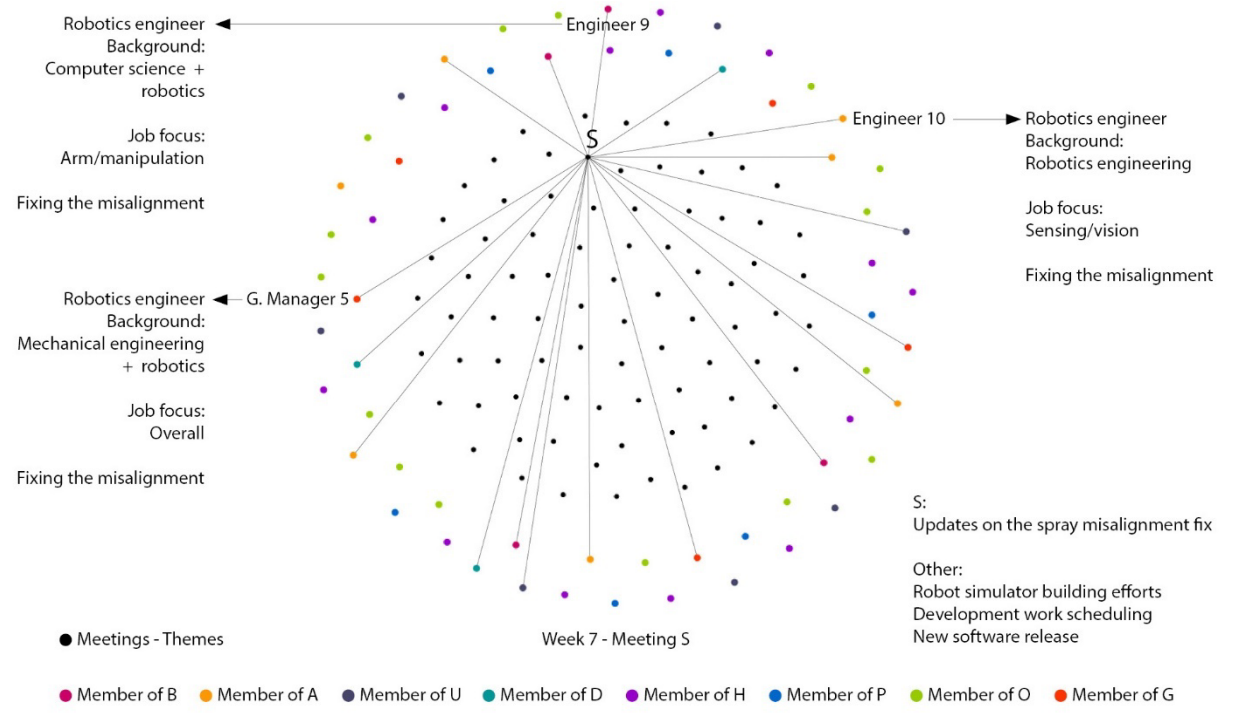


Figure 11. Meeting S showing the diverse technical expertise channeled into fixing a system issue, in this case, spray misalignment before the job site pilot

of the fix, Engineer 9, who specializes in robotic arms and manipulation, had already double-checked the math behind the robot's camera to determine whether the correct frame came back to the software and updated the code accordingly. Then, Engineer 10, who focuses on sensing and vision, continued with the second part, which consisted in looking at the role of camera calibration by deliberately manipulating it to see the extent of the misalignment caused by her manipulation. Once this relationship would be known and the misalignment distance would be fixed, the last step would be to recalibrate the cameras for accurate alignment. A few days later, in Meeting T, Engineer 9 announced that the fix had been implemented and was ready for further testing (Figure 12). From a process standpoint, Engineer 7 and General Manager 5 suggested running tests with multiple machines with different camera calibrations. These tests verified that the fix worked successfully on different machines and the L4TS system was ready for the job site.

The important element demonstrated by the interactions in Meetings F, S, and T was the interdisciplinarity and diversity of expertise informing the development of Basil. In this case, it served to establish the best practices for the newly launched L4TS to maintain the expected finishing quality. These meetings also reveal smaller networks in which issues are traced back to operation, process, or software. Although the final validation of L4TS was sought from the in-house finishers, as discussed in Meeting F, their feedback would later be delivered by Engineer 1 to Team O to inform the team's system productivity test. As Drywall Finisher 4 explained in the meeting, the finishers' brief experience with the system was considered one potential reason behind the minor quality imperfections at the time. This information would later be relayed to Team O, and the team would shift focus to developing the best L4TS user practices to improve quality and train the in-house finishers accordingly. Through Engineer 1, who bridges Team S with the finishers, and Engineer 7 who mediates between Team P and the finishers, the minor quality issues brought up by Drywall Finisher 4 would also be shared with these teams to investigate them on the software and

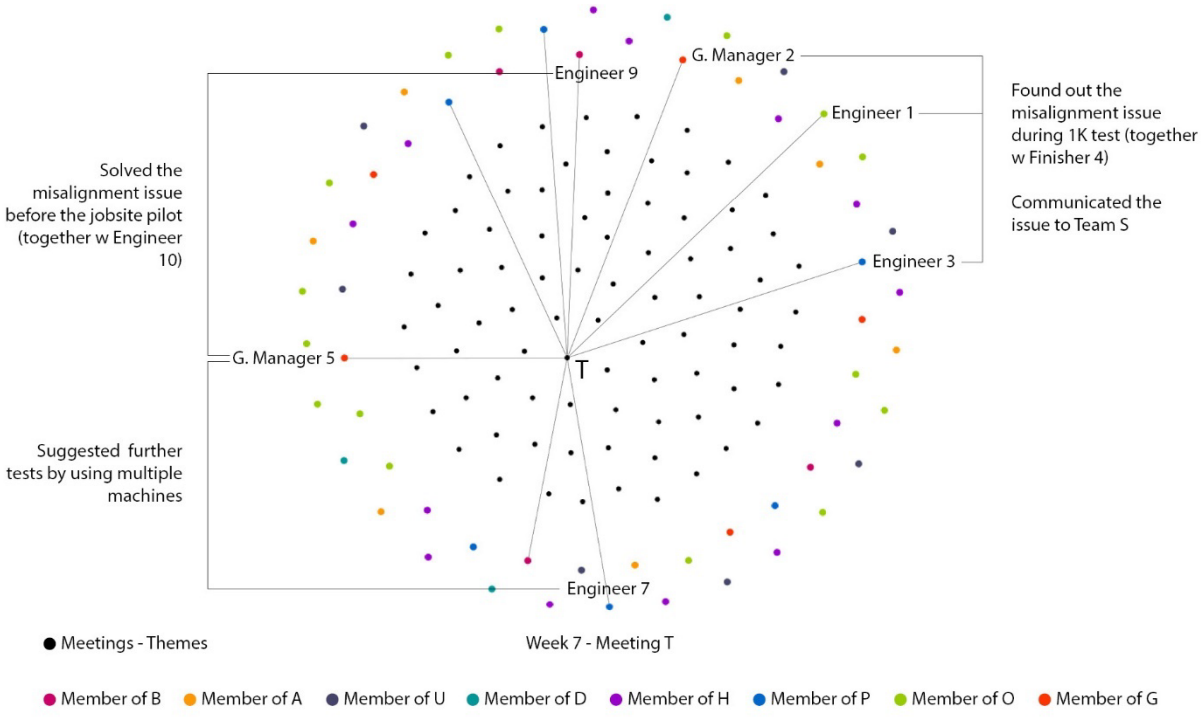


Figure 12. Meeting T showing the network that traced an imperfection in the machine's finishing outcome back to software and hardware

process sides. Similarly, Meeting T demonstrates the network in which an imperfection observed in a test is traced back to the software and hardware through multiple members of Teams O, S, P, and G (Figure 12). In this case, Drywall Finisher 4 and Engineers 1 and 3 found the issue during the 1K test, Engineer 1 and General Manager 2 brought it to the attention of Team S, and Engineers 9 and 10 and General Manager 5 came up with the fix before the machine went to the job site.

Aside from interdisciplinarity and diverse expertise, another notable element I captured within the network is new forms of labor, which are essential to the successful conduct of the robotic workflow and become attached to the drywall finishing trade. One of these is robot operator training, which I initially discovered by surveying the member and theme data for the visualization tool and learned more about later. In its initial years, Canvas operated as a drywall finishing subcontractor, bringing and operating Basil on customers' job sites with the in-house finishers and delivering jobs. Around the time of the study, the company was establishing an additional lease model, in which contractors would rent the machine to operate it with their own finishers. Designed and offered by Training Manager 1, the operator training introduces the machine and the typical workflow I previously described to those finishers and new Canvas employees during a two-day session at the warehouse. The basics of safety, installing the sprayer and sander, positioning the machine, spraying or sanding around wall penetrations, troubleshooting, cleaning the sprayer, and doing blend-ins are among the items on the curriculum that I learned from Training Manager 1 during my own training session and an interview.

Although Training Manager 1 is the central actor in charge, the training is implemented via coordination between him and a network of members, as I initially found out through the visualizations (Figure 13). When I asked about these members and interactions, Training Manager 1 listed "gear, machine, walls, subject matter experts, marketing, and UX" on the whiteboard and said, "[these are the] major resources that need to be aligned to deliver a class." He then explained each member's role in detail. The provision of gear, including side supplies and equipment for the machine, is handled in Team O by Training Manager 1 and Engineer 6. In the same team, Training Manager 1 and Engineers 1, 2, and 5 coordinate the booking of the best-fitting machine and the partitions in the wet testing area. The subject matter experts for the training are the in-house finishers; Drywall Finisher 4 provides a short lecture, and Drywall Finishers 1, 2, and 5, "[the] teacher's assistant[s]," as Training Manager 1 put it, perform hands-on demonstrations of sprayer cleaning and blend-in. The sales and marketing team, together with Engineers 5 and 6, works with the customer to schedule the operator candidate's training and conduct exit interviews with them. Engineer 8 from Team U also interviews the candidate to collect feedback on the UI/UX. The design of the training is also the result of multiple members and interactions. A major effort over several weeks during the study focused on adjusting the existing training for the new L4TS system. The new workflow and the seam review best practices were incorporated into the curriculum.

The other new form of labor that I identified by surveying the data for the visualization tool is site support. "My job is to basically help [finishers operating the machine] troubleshoot issues" is how Site Support Manager 1 later described the scope of this work to me. "My ultimate goal is to keep the machine running on site," he summarized; whenever the workflow is disrupted by an issue, he receives a call from the operator and guides them through how to fix it. Disruptions may be due to a minor system error such as network connection or arm faulting or, often, an operator error, such as mishandling the UI sequence. What Site Support Manager 1 communicates during the support call is then what he characterized as "workarounds" or "tried-and-true solutions" that have been established during previous tests or troubleshooting efforts. An example of a tried-and-true solution that he gave me was the "machine restart," which successfully solved network connection issues

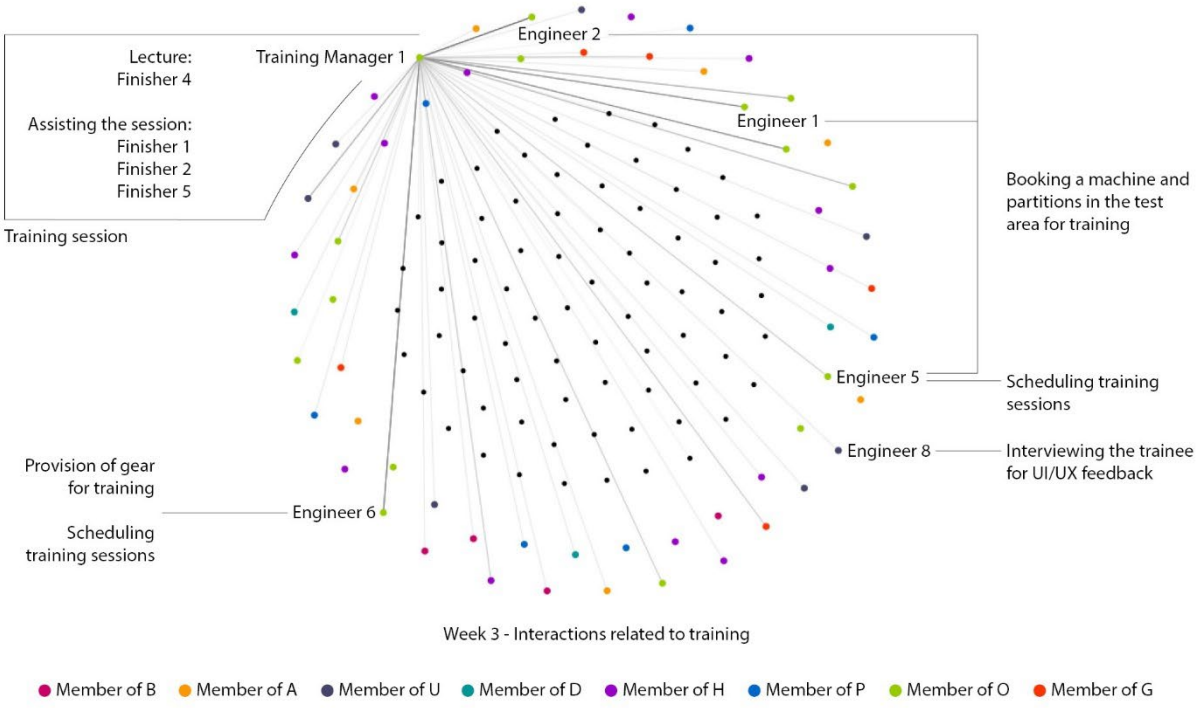


Figure 13. Network of members coordinating with Training Manager 1 and their roles in implementing the robot operator training

between the machine and the controller. Site Support Manager 1 also organizes the recording and communication of issues to relevant teams in an internal pipeline, which is referred to as “ticketing.” In each ticket, he documents how and in what part of the workflow the issue happened and how much downtime it caused. Each ticket is reviewed by Team O, and, if necessary, the issue is handed over to Teams S, H, or P to search for the root cause.

Coordinated between the in-house finishers and members of Teams O, S, H, P, and G, site support work also involves a network of members and interactions I first got a view of on the visualizations (Figure 14). Part of these interactions concerns scheduling, which is carried out between Engineer 5 and Drywall Finisher 4, who arrange the timeline of job site deployments with customers, and Site Support Manager 1, who schedules site support hours accordingly. During job site deployments, Site Support Manager 1 is remotely in touch with the in-house finishers if they are on the job site operating the machine or guiding the customer’s finisher or with the customer’s finisher directly if they are the sole operator. Rarely, when an issue requires an on-site fix, Site Support Manager 1 coordinates with Maintenance Technician 1 for them to go to the job site and fix the issue. Ticket review occurs between Site Support Manager 1 and Engineers 1, 2, and 6 within Team O as well as General Manager 6 from Team G and Engineer 11 from Team S. This helps to “add any missing context that might be helpful later on,” said Site Support Manager 1, meaning that a test instance of an issue that might be brought by these engineers could lead to a better diagnosis. This group deliberates whether the issue should be traced back to the software, hardware, or process. For mid-to-long-term feedback, Site Support Manager 1 is also connected to Team G and General Manager 6, reporting periodical data on issues, as well as to Training Manager 1 to discuss whether any operator errors can be addressed during the training.

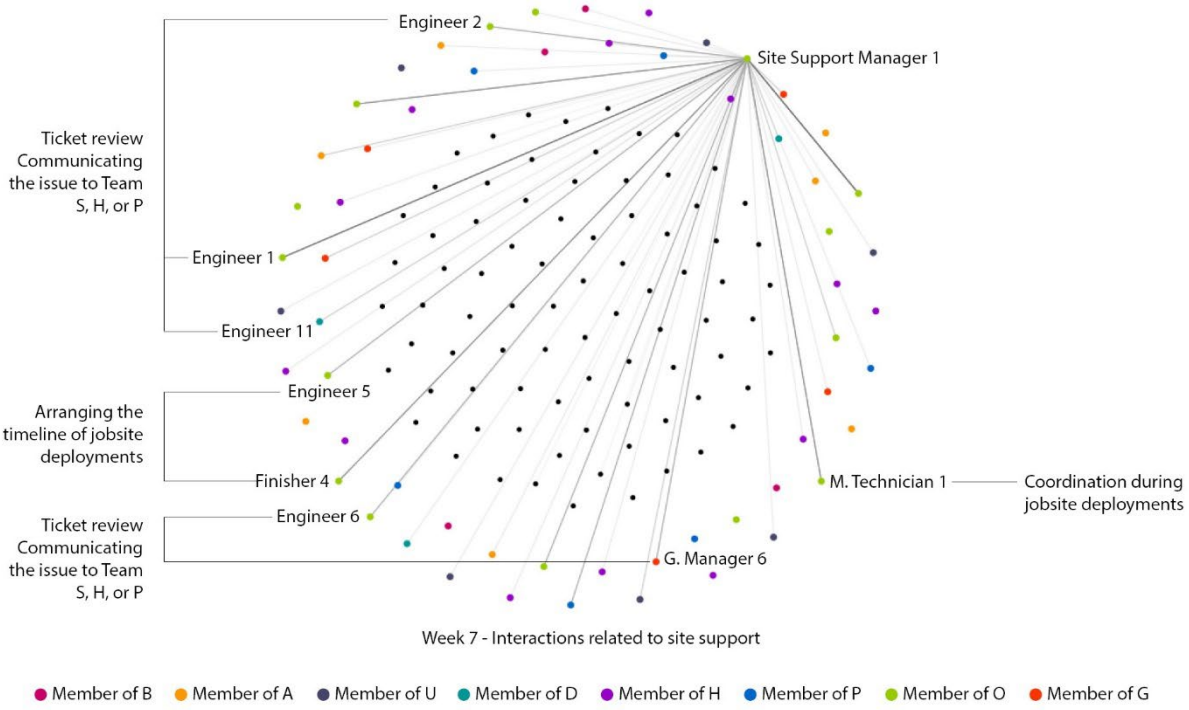


Figure 14. Network of members coordinating with Site Support Manager 1 for site support communication and ticketing of issues

Training and site support illustrate the emergence of new roles and expertise around a new robotic technology in a construction trade. Both are critical to the machine's smooth operation on job sites, and because of this importance, they may become appendage labor inseparable from drywall finishing. Traditionally, union schools are where knowledge of steps and tools in drywall finishing is learned. The robot operator training in its current shape is, in a way, an extension of this traditional entity, teaching a new tool and a new way of performing one step with it. In the near future, the training may also become a distributor of training expertise, as Training Manager 1 foresees a demand among unions or contractors for training and certification of their own in-house trainers. Similarly, site support is an addition to an existing labor group, which Site Support Manager 1 described as "an entire ecosystem of people dedicated to coming out to the site and fixing machines [cranes, backhoes, power tools, etc.]." However, unlike with this kind of conventional equipment, the issues faced when dealing with a more sophisticated and less familiar machine are, in Site Support Manager 1's words, "more nuanced." Therefore, the support is not merely about recalling and servicing the machine and sending it back to the job site but is often a real-time collaborative effort with the finisher to find and implement a fix. "That [form of site support] will become a more common role," said Site Support Manager 1, reflecting on this interactive support emerging with robotic technologies that are new and distinct from previous generations of machines on sites.

A further question related to this body of distinct interdisciplinary interactions and new technical labor I described in this section concerns the kind of change it delineates for drywall finishing as a profession. From a sociological perspective, professions are socially contingent entities, dynamically constituted and reproduced based on their participants' interactions and knowledge.[10] Based on this

10 Based on the "interactionist" perspective (within the sociology of professions—a particular strand in sociology focusing on professions), which foregrounded the role of social constructs and interactions as well as knowledge and skills in the development

framing, one change brought about by the technical labor at Canvas is the addition of new participants who serve the drywall finishing trade without performing construction labor. While they design a technological instrument and integrate it into the trade, this specific group of robotics, software, mechanical, and civil engineers constitute more than a conventional technology supplier and become participants in the trade. Observing their engagement with drywall finishing, which was essential to Basil's success in the lab and on the site, made me see their labor as hardly separate from the trade and these technology design laborers themselves as hardly external to the construction workforce. Aside from the engineers, professionals carrying out new roles that emerge around Basil, including the training and site support I foregrounded, are also among these additional participants. A wider dissemination of machines may produce more of these professionals, situated at the intersection of robotics and the trade in the future.

Another change I view that the technical labor at Canvas signals is the potential emergence of new roles for tradespeople. While engineers become participants in the trade, finishers become participants in the design of a technological instrument for their work. Currently niche and limited, these roles may also propagate with a possible increase in the use of machines. Lastly, interactions between these diverse professionals—shown via the visualization tool—and between their knowledge, which I observed shaping various system components and operations every day, produce new knowledge for the trade. This new knowledge is the outcome of a merging of robotics and construction knowledge worlds, in which it is unclear where one ends and the other one begins. The two worlds serve each other reciprocally. Specifically, this new knowledge produced at Canvas includes an alternative material process that I outlined in the previous chapter, new skills and a new way of performing part of the finishing work that I described previously in this chapter, and know-how in robot integration into the trade that I highlighted in this section.

Lastly, to hear the engineers' and managers' opinions on this convergence of technology design and the drywall finishing trade at Canvas, I asked whether they saw their daily work as part of construction labor. In summary, a consensus and two divergent perspectives emerged. Participants commonly agreed that the development and adoption of a single-task robot involves what General Manager 1 encapsulated as "proximity to the end user," which I call, expanding her definition, "proximity between the technology and the trade it is designed for." Unlike some general-purpose technologies, such as email software, it is difficult to break this proximity with the profession when developing a single-task robot, and it may not be wise for successful adoption anyway. Canvas's initial position as a drywall finishing subcontractor and its transition into a technology provider still closely involved with contractors, trade unions, and finishers at the time of the study is part of building this proximity that contributes to technology development. So is having construction knowledge in-house next to robotics, as I described through several examples.

Aside from this consensus, perspectives on the convergence of technology design and the trade diverged among the participants. Some saw construction robotics, in particular developing a single-task robot, as highly dependent on construction and, consequently, perceived what they do every day as part of construction labor. Engineer 12 was among them: "you can get stuck into your robotics world as much as you want, but in the end, you're still in the construction world," she said to explain that she considers her work to be at the intersection of the two worlds. From a more balanced perspective, General Manager 1 also stated that her work is to operate in both worlds: "as specificity to purpose and proximity to an end user increase, you are still tech, but you are kind of also in another sphere," she remarked. However, she also added that the work Canvas does is

of professions in the 1970s. See an overview of the interactionist perspective in sociologist Keith Macdonald's (1995) chapter "Sociological Analysis of the Professions."

“indirect to construction” in her mind. Others also acknowledged the fact that construction robotics is highly domain-dependent, but they still saw their work as only externally linked to construction. These participants often viewed themselves as mere toolmakers and their job as solving technical problems that are similar across different systems and get integrated into different work domains. They thought of their knowledge in robotics or management as mobile across those different work domains, which, for example, is visible in General Manager 5’s words: “the knowledge about how robots work is transferable, … once you learn something on how to solve a problem, then it can be applicable to multiple things.”

6.3 Summary

In the previous chapter, I showed that two main lines of effort constitute the technology design process that produced Basil: the learning and examination of conventional drywall finishing steps through the lens of human and machine capabilities and, subsequently, the operationalization of the way the mud sticks on drywall to achieve industry-standard finishing quality. Within this technology design process that navigates between the steps, quality, and technology, I identified the engineers’ conceptions of drywall finishing work. The first one approaches it as a combination of craft well suited to human dexterity and improvisation and layers well suited to machine capabilities, such as handling quantity; the second one takes into account the work’s intrinsic reliance on the interplay between action, tool, and material.

These conceptions shaped the development of the machine in a few ways. Guided by the first conception, the engineers departed from unrealistic full automation ambitions towards searching for an appropriate layer in conventional finishing that a machine could fit in and perform. This led to designing a semi-autonomous workflow, which they situated after the taping and nail spotting steps, eliminating box coating in level 4 finishing and both box coating and manual sanding in level 5. Through this conception, the engineers also departed from the idea of emulating the finishers’ skills while building the machine. Instead of finishers’ dexterous skills or movement sequences, which would be expensive to take on without deliberate reduction, the finishing outcome became their focus. This resulted in an alternative material process around the machine that does not exactly follow the conventional coating order. Guided by the second conception, the engineers focused on experimenting with the material process and finishing quality, as opposed to chasing full automation, which could possibly turn into a costly engineering rabbit hole. This helped them to achieve a reliable industry-standard quality with the machine.

These conceptions and processes that they shaped show a technology design work which is strategically positioned at the intersection of robotics and construction, along with the acknowledgement of construction knowledge as an important resource in this work. As I traced when looking at the first line, Drywall Finisher 4’s sharing of her explicit knowledge of the steps of coating and skills enabled the early engineering team to find a step that can feasibly be executed robotically and guided them in configuring the software-hardware that became the foundation for Basil. The idea that combining a base, an arm, an airless sprayer, and a vacuum sander can make a feasible alternative to box coating and sanding, and can create an effect comparable to the one achieved manually was conceived with the help of this knowledge. Similarly, for the second line of effort, Drywall Finisher 4 importing her tacit knowledge of quality and quality evaluation into system development allowed the team to define the specs against which to benchmark the machine’s outcome and, accordingly, fine-tune the software-hardware towards making the outcome reliable.

As a result of these efforts, Basil and the L4TS system reorganize part of drywall finishing into repeated worker–machine actions, as I documented during the 1K test. The division of labor at the job level is restructured around the machine, with finishers handling the pre-machine steps, blend-ins, manual touch-ups, sanding, and final quality evaluation in conventional ways while a finisher and the machine take care of the targeted spraying. With the addition of targeted spraying, the work at the wall level shifts to exchanging cues in a feedback loop. The 1K test also demonstrated that the success of this new worker–machine way of doing drywall finishing is evaluated via the quality of the sprayed and sanded mud, and debugging efforts often start from such quality evaluation and continue by tracing the issues back into the relevant system component. In addition to the iterative cycle in the early stages, this late-stage moment is where we see the finishers' tacit knowledge of quality in use. This knowledge not only played a key role in informing the early-stage development of Basil but also became a measurer of the machine's performance as the process came full circle, as evidenced by the late-stage 1K test.

In this chapter, I addressed the question of how a state-of-the-art robot reconfigures construction work. I pointed out that Basil, as a new technology, reconfigures drywall finishing work in two ways. First, as I showed based on the job site pilot, it does so by rearranging the drywall finishers' skill set, which involves a degree of deskilling and upskilling. On the one hand, the machine eases the entry for less experienced workers into a high-skilled part of finishing. which otherwise requires manual work, dexterity, and more years in the trade. On the other hand, it also demands new skills in addition to the conventional finishing skill set. I described these new set of skills that the system brings in: intellectual skills for settling with the distinct way the machine produces the outcome and how to plan accordingly, and practical skills for successfully running the workflow with both hand tools and the machine, on partitions and their representations across the physical and the software world.

Second, as I highlighted in my analysis of the everyday work life around the robot, the work is reconfigured by introducing a body of distinct interdisciplinary interactions and new forms of labor intertwined with the drywall finishing trade. The technology design labor at Canvas blends diverse professional backgrounds. It comprises new roles and expertise that help facilitate the machine's deployment and long-term adoption on sites. The changes that this labor delineates for drywall finishing are the addition of technology designers as participants contributing to the trade, the emergence of new roles in technology design for tradespeople, and the production of new knowledge for the trade through new interdisciplinary interactions. While some participants agreed on this specific convergence of technology design and the drywall finishing trade, the two are only externally linked in the eyes of others. I see this convergence as a social reproduction of the trade at Canvas and the two chapters as a snapshot of that reproduction.

Chapter 7

Discussion

7.1 Outline

In this final chapter, I first address how conceptions of work that I traced in current practice positioned relative to the origins in the 1980s. I show that recognition of the complex, layered makeup of construction work that is common across the three cases underlie two different ways of approaching work, indicating a shift between the historical and contemporary sites. The first way is the high-level descriptions and classifications of engineers at the DCE concerned with the general character and commonalities of building tasks, and the shift is towards the close, situated engagements with workers and knowledge of tasks in ReAC and Canvas to understand specificities. Then I show two distinctions in understandings of construction work. One is viewing parts of construction work as more mechanical and analogous to manufacturing than it is at the DCE, in contrast to the ReAC group's ambition to understand how the mechanical repetitive activities in construction compare to those in manufacturing. The second distinction is between Canvas's recognition of a task's contingency on material as the key element to comprehend in system development, next to the civil engineers' emphasis on machine autonomy, sensing, or motion at the DCE during the initial years of the field. I discuss that these distinctions are linked to the difference in distance between robotics and construction across the historical and contemporary cases.

Then I look at the conceptions of systems across the cases. First, I examine the single-task robot approach prevalent at the DCE in the 1980s next to Canvas, which put this approach into practice at present. From the 1980s, I revisit Miroslaw Skibniewski's (1986) concept of modularity to overcome a major limitation of single-task robots, which would come with high cost for return of productivity in only one task. Turning to Canvas, I discuss that the development of a modular single-task robot like Basil demands a considerable amount of labor. Further, I point out that persistent concerns around single-task robots since the 1980s, regarding the facilitation that these technologies demand on job sites, are still part of the technology design discussions around Basil. Nonetheless, I also point out that these concerns are contingent on the new job planning skills workers acquire to work with the machine. Second, I contrast ReAC's conception of a support robot with the historically prevalent single-task robot approach. I discuss that ReAC's mobilization of machine learning towards labor support and conception of a robot to perform an ancillary activity displays an alternative to single-task robots designed to handle primary building activities.

Next, I examine how past aspirations to reconfigure construction work in the 1980s unfold in current practice. Based on a meta-review of construction robotics, I first show that early visions of robot-based automation in construction at the DCE in the 1980s remain a driving force of technical research. In this context, I demonstrate that the two contemporary cases (ReAC and Canvas) adopt alternative approaches foregrounding robot support and worker–robot work-sharing for sensible integration of machines into construction rather than automation. Second, I discuss that productivity, which is a principal motivator of on-site construction robots as evidenced by the discourse at the DCE in the 1980s, has a double-sided role today: it functions simultaneously as a key for the industry to accept robots, with Canvas's Basil ticking this box, and a barrier to imagining alternative systems that do not prioritize it, such as the tool and hardware delivery robot developed in the ReAC project. Lastly, I show that deskilling and upskilling as well as an emerging technical labor at the intersection of construction and robotics constitute the current reconfiguration of construction work around robots. This differs from the early labor displacement projections by the DCE's publications.

Drawing from these findings**,** I finally present the dissertation's three main arguments (that I outlined in Chapter 1), along with a series of reflections that it produced. The first series concerns technology design in construction robotics. Here, I propose approaches and methods for pursuing a sensible

design and integration of robotic technologies into construction as a response to four key issues I identified within the perspectives and practices of workers and technology designers: the situated nature of tasks; skills, deskilling, and upskilling; the fitness of robots and sites; and roles for robots. The second series of reflections is related to pedagogy in architecture, construction, and robotics. I discuss why it is important to incorporate into the pedagogies of these fields concepts and methods for studying, understanding, and articulating new questions about technology design, technologies, and technologies' effects on work. As an example case, I describe my experience of this in a seminar course I offered at MSD-RAS, a one-year post-professional program focusing on robotic fabrication, design computation, and material exploration at UPenn's Weitzman School of Design. The final series of reflections are on the autoethnographic approach in the ReAC study and the way the three customized research tools in this dissertation are positioned within and serve the ethnographies of ReAC and Canvas. I conclude this final chapter with a brief note on the dissertation and its relevance.

7.2 Between Past and Current Conceptions of Construction Work

7.2.1 From High-Level Understandings to Situated Engagements With Construction Work

As I discussed in Chapter 2, the discourse and instruments (Warszawski and Sangrey's (1985) taxonomy and Skibniewski's (1986) modular design framework) at the DCE manifest the engineers' conception of construction work as a complex makeup of site conditions and activities. The engineers already recognized that this complex makeup poses challenges beyond those associated with developing and adopting robotic systems in manufacturing. With this recognition, the idea that engaging with construction work is a prerequisite for realizing the ambition of construction robots already existed in the early days of the field. This early recognition of the relevance of involving construction at the Department seems to be the main link between the past and current practice. As I revealed in Chapters 3 to 6, a similar conception of construction work as a richly layered practice underlies ReAC's ethnographic and Canvas's drywall finishing studies. Researchers and engineers at these contemporary sites have the same awareness as their predecessors at the DCE in the 1980s: one cannot do much in robotics without understanding the building tasks themselves.

However, there are differences in how researchers and engineers approached work based on this conception in the past and present. As I discussed in Chapter 2, the engineers at the DCE featured high-level descriptions and classifications of tasks to make a case for robots in construction. These descriptions and analyses sought to capture the tasks' general character and commonalities, and overlooked their situated nature and, relatedly, their idiosyncrasies. They represent a one-sided engagement, in which engineers remotely learned and utilized construction knowledge. In contrast, the two contemporary practices have built two-sided engagements including sites, workers, and teams that combine construction and robotics knowledge as participants in the design of systems.

Through the ethnographic study, the ReAC group strived to understand the specificities of tasks rather than their general character and detailed the commonly recognized attributes of construction as I described in Chapters 3 and 4.[1] Further, through the study, we worked to understand the idiosyncrasies of tasks rather than collapsing them into common categories. Instead of descriptions or classifications that abstract away the details and say what a task is, the study produced

1 For instance, the study helped reconceptualize the construction workplace, which is briefly characterized as unstructured in publications in the 1980s, as not unorderly and loosely structured through the organization of zones. It captured the different extent to which each task or activity embodies the general qualities associated with construction work, such as physicality, repetition, and mobility.

descriptions that preserve the details and tell how and in what specific site conditions a task is enacted. For instance, the taxonomy approach at the DCE was about categorizing tasks to match them with robot configurations, which ultimately implied that every task could, in theory, be automated with the existing archetypal technologies. Based on a detailed understanding of each task's qualities, in contrast, ReAC's bottom-up search asked which ones are more amenable to robotic support.[2] Canvas also aimed to understand the specificities of tasks. As I described in Chapter 5, the engineering team learned about drywall finishing from the in-house tradespeople, including the steps of this task and finishing quality, in the company's early days. Similar to the taxonomy approach, the engineers assessed the steps to find one for which a machine could realistically be designed. Departing from this same approach, however, they selectively searched for one that was also appropriate to automate.

There are also distinctions that I traced in understandings of construction work across the cases. As I noted in Chapter 2, some engineers at the DCE considered part of construction work, in particular surface finishing tasks, as more mechanical and analogous to manufacturing tasks to make a case for robotization. The nuance that was neglected is that manufacturing work broadly comprises more tightly planned, granular, and repeated cycles of action than on construction sites, which makes it more accommodating to robots, as I described in Chapter 2. Conversely, work on construction sites occurs in relatively more loosely planned and broader repeated cycles of a larger mix of actions and improvisation, often driven more by workers than machinery. The surface tasks demand a different type of job planning and coordination with other trades, along with different quality standards than the tasks that resemble them in manufacturing. In the ReAC group, we started from the premise that manufacturing work and construction work are distinct from each other, but by developing a clearer sense of how repetition is performed in each building task, we gained a better understanding of the nuance between the repetitive steps of construction and those of manufacturing.

The other distinction concerns the recognition of tasks' material-oriented character. Acknowledging the drywall finishing's reliance on the coating material, the Canvas team showed that successful robot development and integration lies equally in understanding the material process and outcome as it does in developing system capabilities. This conception, which pushed the Canvas engineers to start by exploring the way an arm works with the mud and drywall—instead of seeking innovation in computing—and to even strategically compromise certain capabilities, such as autonomy, was mostly missing at the DCE during the initial years of the field. Both Warszawski and Sangrey (1985) and Skibniewski (1986) featured a similar class of tasks that involve coating a surface with substance, such as painting, plastering, and applying mortar, or entail removing a substance from a surface, such as sandblasting, troweling, and brushing. However, their main emphasis was more on advancing autonomy, sensing, and motion to successfully fit machines into these tasks, and less on developing processes to adjust the fitness of machines to reliably work with the tasks' contingencies with materials.[3] Other engineers at the DCE such as Oppenheim, Hendrickson, and Whittaker who developed systems solely worked to advance system capabilities.

---

2 For instance, hanger installation and drywall panel installation were simply grouped in the "attaching" category in Warszawski and Sangrey's (1985) taxonomy, whereas how and in what site conditions attaching is performed in each activity was specified and qualitatively described in the ethnographic study. Accordingly, each activity was assessed to see whether it would be a credible form of support to workers (if it was performed by a robot).

3 For instance, Miroslaw Skibniewski (1986) envisioned quality control with a sensing system as part of the conceptual sandblasting system in his dissertation. Whereas quality control is at the core of system development at Canvas, yet also one of those steps in finishing Canvas did not try to address with machine intelligence. Instead, it is the workers who evaluate the robot's spraying or sanding performance. As further evidence, William Whittaker's (1986) aforementioned proceeding "Construction Robotics: A Perspective" at the third International Symposium on Automation and Robotics in Construction (ISARC) offers a critical point on the field's emphasis in the 1980s. In this proceeding, Whittaker (1986) wrote: "Perhaps the early American views of construction robotics overestimated the need for sensing, artificial intelligence, and autonomy. We speculate that these cognitive attributes will

The main reason I see behind these distinctions is the difference in distance between technology design and construction, which anthropologist Lucy Suchman (2002a) describes as the "distance from the specific sites of technologies-in-use." At the DCE, the worlds of robotics and construction were viewed as distant as I noted in Chapter 2: technologies produced in one world would inevitably disrupt labor in the other. Analyzing construction work and building robots were separate lines of effort. Engagement with construction was remote and one-sided. The contemporary cases show that the distance between the two worlds, which was greater in the 1980s, might be starting to close. In ReAC and Canvas, there is an acknowledgment that robotics is partly dependent on construction. As I discussed in Chapters 4 and 5, the granular view of work beyond generic descriptions in ReAC, and the importance of material and finishing quality in system development in Canvas were shaped by the knowledge of workers. Analyzing construction work and building robots are beginning to converge in system development. Engagement with construction is situated and two-sided.

However, as much as the interdisciplinarity arrangements in ReAC and Canvas facilitate a knowledge exchange leading to more sensible systems and integration, they also involve critical challenges. Because technology design and construction are inherently different disciplinary cultures, even though they are brought closer, it is difficult for the members to adopt multiple positionalities across construction and robotics or fully align their practice with these fields. Most members cannot help but view their identities as well as concepts or problems through the lens of one knowledge world more than the other. I observed this challenge in some Canvas engineers who saw themselves as mere toolmakers free of any entanglement with the domain of construction as I highlighted in Chapter 6. This challenge was also visible in the tension and negotiation between conceptual and technical angles among ReAC members as I discussed in Chapter 4.

7.2.2 Machine Learning for Support: An Alternative to Single-Task Robot Approach

Influenced by the first strand of single-task robot development in Japan in the late 1970s and early 1980s, the DCE had these technologies as main reference.[4] The CECRL's RRV and Terregator were mobile bases designed to turn into various single-task robots with appropriate tooling. Both Warszawski and Sangrey's (1985) taxonomy and Skibniewski's (1986) modular design framework were guides for the development of this type of robots. Their common purpose was to expand these machines' utility with minimum changes in configuration. This was meant to help overcome a major limitation of these technologies, which contemporary researchers such as Fopefoluwa Bademosi and Raja Issa (2021) still point out today: the high cost of developing or acquiring these technologies for a potential productivity gain in a single task.[5] Through their taxonomy, Warszawski and Sangrey (1985) posited that all ten activities that constitute all major building tasks could be covered with four robot configurations with various tooling options, as I mentioned in Chapter 2. Through the modular design framework, Skibniewski (1986) argued that similarities in tasks and in the design of robots could be leveraged to build modular systems to perform a set of similar tasks. Both guides aimed to show that the high initial cost might be distributed across multiple tasks.

Canvas's Basil is a present-day addition to the single-robot strand in construction robotics. With its ability to switch from spraying to sanding with only an end effector change, this machine is, in a way, an implementation of the modularity that Skibniewski (1986) envisioned earlier. However, the Canvas case shows that the labor needed to facilitate even a tiny scale of modularity to expand the

eventually dominate construction robotics. Nonetheless, teleoperators and programmed machines have both short- and long-term relevance."

4 This class of robots is what the civil engineers denoted with the word "construction robots."

5 Other studies, such as the ones by Mi Pan and Wei Pan (2020a), Juan Manuel Davila Delgado and coauthors (2019), and Thomas Bock and Thomas Linner (2016) have also pointed out the issue of high costs related to implementation.

utility of a single-task robot is considerable. Developing a process, adjusting the software, conducting rounds of tests, and carrying out much internal communication are parts of that labor as I described in Chapters 5 and 6.

The Canvas case also shows that the practical concerns, which arose with the initial single-task robot experiments in Japan, are still relevant and discussed decades later. For instance, the concurrent operation of a single-task robot and a worker crew performing two different tasks in the same zone is not always possible due to safety and productivity reasons (Bock & Linner, 2016). Finishers and engineers at Canvas also pointed out this challenge, and relatedly the importance of task scheduling around a single-task robot. The amount of time and skill that must be invested in the "setup and programming" of these machines on site is the other potential issue that has been under discussion (Bock & Linner, 2016). Helping to reduce that amount is both the training and the site support managers' job at Canvas. Nonetheless, by closely examining the workflow with Basil, I observed that these concerns are not default facts but are instead contingent on foremen's and machine operators' planning skills, which I described in Chapter 6—more specifically, knowing how to plan the conventional and machine steps of the task and the space needed for better integration of the machine on a site.

Next to the long lineage of single-task machines that constitute most of what we know as construction robotics, ReAC's tool and hardware delivery robot demonstrates an alternative conception. A major portion of single-task robots conceptualized or built over the past four decades are designed to perform primary building activities, such as drilling, fastening, spraying, or installation (Bock & Linner, 2016).[6] While some of these systems are designed to free workers from heavy construction tasks that are currently conducted with operated equipment, some are designed to take over parts of tasks that workers may want to keep. The integration of these robots may demand a transition from workers to machines in certain activities. It may require upskilling or shifting part of workers' roles to operators or supervisors, as I also discussed in the case of Basil in Chapter 6. In contrast, ReAC's tool and hardware delivery robot leverages machine learning for performing an ancillary—rather than primary—activity to support labor.[7] As I discussed in Chapter 4, this preserves the primary activities that the carpentry workers invest time in to train for, use skills to enact, and value. The use of this robot also requires a set of new skills, such as the ability to send and call it back via a controller. However rather than the carpentry workers adapting to carry out the formwork workflow with the system, the system itself is designed to be an adaptive part of the workflow—without requiring any change in workers' activities.

7.3 Between Past Aspirations and the Current Reconfiguration

7.3.1 From Visions of Automation to Robot Support and Worker–Robot Work-Sharing

"Is the talk of automation overblown?" asks sociologist Aaron Benanav (2020) to question the gap between the automation discourse in today's technological and political circles around AI and robotic

6 See Bock and Linner's (2016) survey of single-task robots, which comprises two hundred examples. The majority of these have been developed and deployed in Japan, where the idea of single-task robots emerged.

7 The tool and hardware delivery robot is close to the category that Bock and Linner (2016) identify as "site logistics robots," which includes mobile large-scale material delivery systems developed by Japanese contractors. However, there are a few fundamental differences. These Japanese single-task robots are designed to serve as one node in an automated network of vertical and horizontal delivery systems transferring large-scale materials (Bock & Linner, 2016). Relatedly, these robots move along designated routes that are marked and pass materials to the next system at designated points (Bock & Linner, 2016). In this sense, they are part of a setup that requires a partial structuring of sites. Conversely, ReAC's tool and hardware delivery robot is a learning-based system designed to operate with workers as part of a specific task, as opposed to working as a connector on a standalone logistics line feeding various tasks. Basil can adjust its moves in real time according to workers' movements and obstacles. With this capability, it does not require any significant site adjustments.

technologies and the speculated implications of these technologies for labor. Benanav (2020) maps "four principal propositions" on which the contemporary automation discourse is built, two of which are also relevant to the 1960s–1980s technology-labor context I briefly navigated in this dissertation. The first is that the discourse claims there is ongoing "technological unemployment" and, relatedly, projects "a largely automated society" that has long been dreamed about and feared.

Advancements in robotics from the 1960s through the 1980s triggered excitement and high hopes for robot-based automation in manufacturing among industrialists and researchers, including those at the Robotics Institute, as I described in Chapter 2. Similar excitement and hope for a similar transition in construction pervaded the DCE. In the 1980s, while robot-based automation in manufacturing was a reality in progress, robot-based automation in construction was only a vision, which was becoming clearer. Therefore, while technological unemployment was partially a current reality and partially a future concern in manufacturing, it was solely the latter in construction. Among the historical documents I analyzed in Chapter 2, Ayres and Miller's (1981) report provided estimations of worker replacement based on a survey with members of the Robot Institute of America, listed the states that might be affected the most, and even discussed how labor unions could be involved in this forthcoming transition. Meanwhile, citing one of Ayres and Miller's (1983) later assessments, Warszawski and Sangrey's (1985) paper predicted a similar trajectory for construction, despite the fact that robots were not really present on job sites in the United States.

Still, as I revealed regarding the aspirations of the 1980s at the DCE, construction was not excluded from the projections of a largely automated society at the time. A major aspiration of engineers around robots was to improve the declining productivity in construction, which represented a significant share of the US economy, as I traced in the publications of Whittaker (1986), Warszawski and Sangrey (1985), and Skibniewski (1986). To streamline the work towards productivity, the other aspiration consisted in structuring the on-site workflow by turning it into a software operation and the worker into the overseer of an autonomous robot fleet, as I showed with Whittaker and Bandari's (1986) multi-robot system. The discourse within these particular academic circles of the past, like the one Benanav (2020) pointed out within the technological and political circles of today, was going faster than the technologies or their real-world implications on sites. The engineers painted a bold vision of automation in construction.

Thus, I reiterate Benanav's (2020) question for the past and ground it in the context of construction: was the talk of automation in construction overblown? How do the aspirations I traced within the DCE in the 1980s, which played a major role in kick-starting on-site construction robotics in the country, unfold in current practice?

Reviews of the field since the 1980s seek to answer this by focusing on technological progress. They survey advancements in labs and applications on sites and, in a way, assess where progress towards the initial promise stands. For instance, around the tenth anniversary of the first ISARC at CMU, a survey of the conference's proceedings by Luis Ochoa-Franco and colleagues (1994) showed an upward trend in the percentage of commercial-grade systems from 1990 to 1993, coupled with a greater number of implemented systems than concepts in 1993. "Construction automation [was] maturing," according to the researchers (1994). Still, "much remains to be accomplished," Miroslaw Skibniewski (1996) concluded based on his survey of developments covering ten academic and industry research centers in the United States. Acknowledging the "limited capabilities" of systems at the time, the researcher (1996) called attention to "the redesign of traditional work tasks" to accommodate them. The situation on sites also substantiated his point. A broader survey by Abraham Warszawski and Ronie Navon (1998) revealed that only twenty-six out

of seventy-six systems globally (in eight countries, including the United States) were used on sites, with only five of these deployed more than fifty times. These systems were mostly teleoperated and far from realizing the vision of fully autonomous construction (Warszawski & Navon, 1998).

From this same technological progression angle, the first two decades of the 2000s have not been held in higher regard. According to Hyojoo Son and colleagues' (2010) review, construction robotics publications, which constituted seventy-one percent of ISARC proceedings between 1990 and 1992, dropped to less than half between 2006 and 2008. Applications were still limited to "teleoperated and programmable" systems rather than "autonomous or semi-autonomous" ones, as Kamel Saidi and colleagues (2016) pointed out. More recently, Nathan Melenbrink and colleagues (2020) suggested that there is "very little compelling evidence" on the scale of robot integration in the architecture engineering construction (AEC) industry estimated by "economists and investors." Surveying several on-site systems for earthmoving, excavation, pile driving, structural assembly, additive manufacturing, and finishing, the researchers (2020) reported that these systems currently have a smaller effect on automation because only a small fraction of them can "autonomously execute" a task without any worker assistance. Sofiat Abioye and colleagues (2021) identified "costs," "potential job loss," and the "unstructured" nature of sites as obstacles to the strengthening of this effect. Conversely, for Yue Pan and Limao Zhang (2021), technological limitations are to blame, so "equipping the robot systems with more powerful abilities" could scale up their use.

These reviews approach the systems of the past three decades as a monolithic effort on a continuum towards full automation. The implicit focus in these reviews is on checking every few years where the state of the art stands relative to the bar set by the vision of the 1980s—that is, how close or far the current state of automation is, what the challenges and limitations are, and what the next steps could be to reach this goal. Unlike these bird's eye reviews, in this dissertation, I zoomed in on and closely engaged with 1980s CMU along with two contemporary practices, ReAC and Canvas. Based on these cases, my answer to the question of how the early aspirations of the field unfold in current practice emphasizes a divergence from those aspirations. While ReAC's and Canvas's systems can be viewed as part of the continuum drawn by other researchers, they also represent a conceptual shift from the bold vision of automation of the 1980s to a more realistic notion of automation built on robot support and worker–robot work-sharing, respectively.

For instance, Melenbrink and colleagues (2020) outlined the ongoing "trends and developments" in the pursuit of automation. One is "material-robot systems," in which materials or building techniques are redesigned contingently with a robot. Another one is "embedded sensing in building components" to enable distributed robot fleets to work on assigned aspects of building construction. The other lines of development concern systems for "auxiliary construction tasks" to build the temporary structures used in construction, systems for "site preparation and substructure tasks" to handle earthmoving, excavation, or pile driving, and coordinated multi-robot systems for "task integration."[8] In another recent review, Bo Xiao and colleagues (2022) divided the developments into more general categories: "tasks," including on-site operations, off-site manufacturing, and additive manufacturing, "algorithms" for sensing and autonomy, and "collaborations" in human–robot and multi-robot configurations. Other than the self-evident categories such as "on-site operations" or "algorithms," the ReAC and Canvas cases do not necessarily fall into one of these.

Informed by an ethnographic study of construction, ReAC's technology design process foregrounded ancillary activities in tasks and ended up focusing on tool and hardware delivery. Despite being

8 In addition to the ones in the text, another "ongoing trend and development" that Melenbrink and colleagues (2020) listed in their review is "interdisciplinary fundamental research."

situated in carpentry formwork, which Melenbrink and colleagues (2020) identify as an auxiliary task, the final system is clearly different from the lab-based assembly systems they describe under this line of development. Rather than take over the entire formwork assembly, the ReAC's system is designed to perform one of many activities to support carpenters while they assemble formwork panels, and it is much closer to operating in an actual workflow. The system is also different from those that Xiao and colleagues (2022) classify as human–robot collaborations. These systems are designed to shift workers to off-site control through teleoperation or motion tracking, which is far from being accomplished in the current organization of sites. Instead of facilitating the remote operation of a task, ReAC's system is designed to support workers performing it on site.

On the other side, driven by a re-exploration of drywall finishing steps and material, Canvas's technology design process represents an approach similar to what Melenbrink and colleagues (2020) highlight in lab-based material-robot systems performing assembly and fabrication with experimental materials. However, rather than innovating a material-robot system not yet applicable and foreign to the existing construction industry, Canvas focused on drywall finishing as an existing material process and modified its steps while designing a system to deliver industry-standard quality. Again, unlike Xiao and colleagues' (2022) human–robot collaboration, the final system is designed to semi-autonomously perform a particular step in drywall finishing, where the task becomes shared between a worker and a robot.

With this conceptual shift towards robot support and worker–robot work-sharing, the ReAC and Canvas cases point to the emergence of a more grounded take on automation, which acknowledges the complexity of tasks and sites as well as the limitations of robots in current practice. As an alternative to pushing the state of the art forward towards automation, which has arguably not been a very fertile strategy over the past four decades, both groups work practically with the state of the art and what exists on sites. In the ReAC project, this is built into the project's conceptual framing, which approaches the issue of technology design through a sociotechnical lens. At Canvas, it stems from a practical approach aimed at building a commercial-grade machine, which has to align with and survive the immediate realities of a task, sites, and the industry to be accepted.

7.3.2 Productivity as a Key and a Barrier

As I mentioned in the previous section, a major aspiration underlying the work at the DCE was to improve the declining productivity in construction. The construction industry has continuously been a major part of the US economy, contributing an average of four point three percent to the US gross domestic product between 1950 and 2020; however, the decline in the sector's productivity that the civil engineers discussed at CMU in the 1980s has not been overturned since then (Goolsbee & Syverson, 2023). Therefore, almost four decades later, productivity is still an unchallenged variable in construction robotics. Productivity improvement remains the main promise of the field of construction robotics to the construction industry. It is the biggest card that researchers and engineers have to play when pitching new systems and one of the success criteria the industry expects these systems to fulfill and evaluates them against. In this sense, productivity is the key to opening the gates of the industry to robotic technologies.

The significance of productivity as a stimulus in the two contemporary technology design cases I looked at differs from what prevailed in the 1980s as well as from each other. Situated within a university, the ReAC group worked with sponsors and industry partners but did not have to deal with any commercial or sales pressure. Productivity was an indirect consideration in the sense that successful robot support could contribute to it, but it was not the core focus or a major part of our

conversations. Supporting labor was the main motivation. Not being exposed to such pressure perhaps allowed us to patiently study multiple tasks in context, focus on ancillary activities that might not have direct measurable effects on job progress or productivity, and eventually develop a system that performs one of those. Although the system may have indirect effects on productivity, which we did not assess, reducing the labor of a multi-tasking carpenter was the benefit we put forward to make a case for it. I doubt that this alone, without proof of any productivity gain, would be convincing if we were to pitch this system to a contractor. However, the absence of this requirement allowed us to conceive and demonstrate an alternative role for a robot in construction.

At Canvas, productivity is also not the main aspiration behind Basil; yet, as the publications of the 1980s reflect, it is the principal criterion on which the company must evaluate Basil to promote it to the industry and the principal condition for industry actors to accept it. In this sense, Basil is not free from what I call "productivity pressure." Because the system has been proven to match industry-standard finishing quality in a quantity that is deemed productive, it is accepted by the contractors or job owners that have so far used it. However, this requires a considerable amount of effort from the engineers and adjustments from the workers who use it, which I described in Chapters 5 and 6. Along with the finishing quality, Canvas uses the productivity of the robotic workflow to evaluate their system, as I showed with the 1K test. Some of the new skills I described, such as planning the sequence between conventional and machine steps and keeping progress on manually and machine-coated parts at the same pace, are also geared towards maintaining productivity.

ReAC's support robot prototype and Canvas's commercial-grade single-task robot are different types of systems developed on two very different sites and with different motivations. Nevertheless, the difference between having to live with productivity pressure and being liberated from it points to an interesting potential. That is, temporarily divorcing technology design from productivity may open up space for imagining other systems with alternative roles, some of which may also end up serving the decades-long goal of improving productivity.

7.3.3 Emerging Technical Labor in Construction

The notion of robot-based automation is often accompanied by one popular question: will robots take our jobs? This long-standing, often speculative inquiry into the future assumes that job displacement is an inevitable consequence of automation in construction—colloquially, a matter not of "if" but "when." Presenting job displacement as destiny obscures the possibility of designing more human-centered technologies and technology integration. It gives a wrong impression about the actual capabilities of current systems. Focusing on the sociotechnical effects of future robots diverts our attention from understanding the current effects of existing ones, which are more nuanced than job displacement. Through the cases in this dissertation, I had the chance to examine the aspirations of the past and the reconfiguration of a specific construction task around a commercial-grade robot in the present.

Technological unemployment in construction was discussed as a potential future concern in the 1980s in the US because there were advancements but no significant sign of robots on sites.[9] While Warszawski and Sangrey's (1985) paper anticipated "[labor] displacement" and a "redefining [of] the task[s]," which may cause issues of "labor cooperation" in robot integration, Skibniewski's (1986) dissertation evaluated the "cost and benefits" of "[worker] replacement." Today, technological

9 For instance, in "Robotic Applications to Construction" referenced earlier in the dissertation, Miroslaw Skibniewski and Jeffery Russell (1989) surveyed only thirteen prototypes developed in the United States by 1989. The majority of these were developed in universities.

unemployment in construction is still discussed as a potential future concern partially because the technologies are not fully mature. The post-2010 reviews of robotics in construction I touched upon in the previous section imply that neither robots nor sites are ready for full automation. Additionally, the integration of robots into jobs is still not substantial. Despite a considerable increase since 2017, Geno Armstrong and colleagues' (2023) survey of one hundred and twenty-one engineering and construction companies worldwide (including a sample from the United States) revealed that only seven percent of these reported using robots in all their projects, whereas thirty-six percent "started [using robots] with a few projects."[10] Meanwhile, workers do not really expect any immediate changes in terms of robot adoption despite experiencing some amount of anxiety about potential job losses, according to YeSeul Kim and colleagues' (2022) interviews.[11]

Knowing the true complexities of their work, the New Scaife Hall workers also do not expect any drastic changes on sites any time soon. A more realistic effect of the dissemination of robots they foresee is having to work with new instruments as well as potential deskilling and upskilling, which the small group of finishers at Canvas are already open to exploring in another part of the country. As I described in Chapters 5 and 6, these finishers' work is reshaped but not exactly into the clean division of labor Whittaker and Bandari (1986) aspired to in the 1980s, in which robots autonomously work and a worker oversees a full software operation. Instead, it is reshaped into a division in which the steps constituting a task are shared between a worker and a semi-autonomous robot: the worker handles the initial steps, uses the robot to semi-autonomously execute the main step in a software operation, oversees the execution and the result it produces, and handles the subsequent steps. Instead of rendering trade skills obsolete, the robotic workflow is designed to be accompanied by those skills in the steps for which the machine is not designed, along with those in quality evaluation and manual fixes. The system creates a degree of deskilling as it levels the playing field for less experienced finishers for part of the task. Nevertheless, skilled finishers are essential for accomplishing a job with the robot. The system also requires a degree of upskilling by demanding intellectual and practical skills in addition to conventional ones.

From the 1980s to the present, the state of the art in construction robotics has not generated the social cost for labor anticipated by the DCE's figures. The "automation-driven labor share decline" in construction between 1987 and 2016 amounted to ten percent, much lower than in other industries, such as mining, automotive, electronics, or food manufacturing, as economists Daron Acemoglu and Pascual Restrepo's (2022) study shows. Perhaps due to the distance between the lab and the site at the time, the engineers' discourse at the DCE portrayed technology design and construction as two separate worlds: the former as the world of invention and the latter as the world that is unavoidably affected. The robotics workforce that produces knowledge and technologies and the construction workforce whose work technologies could disrupt seem to be viewed as two separate entities. Canvas, as a prominent actor in the industry, shows that these two separate worlds and entities not only got closer but are becoming intertwined in current practice. The engineering team at Canvas is composed of members from these two workforces. As I discussed and emphasized in Chapter 6, this team brings a new set of professionals who are not the conventional AEC actors, their daily work interactions around the system, and new roles and expertise such as robot training and site support into construction. All these constitute an emerging technical labor in the social reproduction of the

10 Geno Armstrong and colleagues (2023) conducted this survey under the global consulting firm KPMG. The survey involved two hundred and sixty-seven participants from engineering & construction companies (one hundred and twenty-one) and project owners (one hundred and forty-six) around the world. According to the survey, the current forty-three percent robot use among engineering & construction companies (seven percent using in all projects plus thirty-six percent starting) was ten percent in 2017. In a similar survey, the robotics company ABB (2021) reported fifty-five percent of one thousand nine hundred decision makers in the construction industry from the United States, Canada, China along with seven European countries stated robot use.

11 Kim and coauthors (2022) interviewed seventeen workers from various trades.

profession. Thus, the Canvas case displays a less visible reconfiguration around a robot, which differs from the commonly speculated job displacement and may be further explored in future research.

7.4 Final Reflections

As I pointed out in the introduction, as much as work is a form of performance that robotic systems are developed to accomplish, it is also a less visible knowledge source, which engineers and researchers have sought to understand in developing and informing these systems. I discussed that engineers and researchers build conceptions of construction work and that these conceptions are reflected in the ways they study, analyze, and understand it. Based on that premise, I explored these conceptions and the linkages and shifts between them in history and in current practice by closely looking into three prominent sites of robotics—CMU-RI and DCE, ReAC, and Canvas—in this dissertation.

Overall, the cases demonstrate that technology designers' conceptions of work, alongside technical artifacts, decisively shape system development in construction robotics. The exploration of conceptions shows a linkage, which is the recognition of the complex, layered makeup of construction work common between the cases, and a shift in the way in understanding this complex entity: from high-level descriptions and classifications of building tasks concerned with their general character and commonalities at the DCE to close, situated engagements with workers and knowledge of tasks aimed at understanding their specificities at ReAC and Canvas. Each case also reveals different conceptions, such as ReAC's understanding of nuances in tasks' qualities and Canvas's understanding of the tasks' contingency on materials, underlying different technology design directions that shaped different systems.

The cases also show that the participation of designers and workers in each other's knowledge worlds helps to conceive systems with alternative roles (compared to that of single-task robots) as I showed in ReAC, and to practically integrate robotic technologies into existing workflows as I showed in Canvas. It is important to remind ourselves that capturing knowledge of work in technology design has involved pitfalls in the past as anthropologist Diana Forsythe (2001) has shown in the context of AI. Forsythe (2001) has observed that, while obtaining knowledge of their experts' work, AI researchers overlook the "distinction" between the practice of that work that one could observe in context and understanding of the same work that is described to them by experts. What helped to avert that in the contemporary cases in this dissertation is the two-sided engagement with construction, in which the designers' incorporation of the knowledge of work into their processes is through not only workers' descriptions but also observation and participation of workers. This helped draw the understandings of labor and sites closer to the labor and sites in reality.

Aside from being a knowledge source, construction work is a sociotechnical operation that engineers and researchers have aspired to reconfigure with systems in the past and realize through systems deployed on sites in current practice. I described these as discursive and technical reproduction of the arrangement of workers, instruments, interactions, or organizational structures that facilitate the performance of construction work. Thus, I traced these aspirations by zooming in on discursive and technical instruments of the past at the DCE, and explored how they have unfolded at present by examining the technology design processes and systems of ReAC and Canvas. I showed the current reconfiguration of a building task around a commercial-grade robot at Canvas.

In this exploration, the historical and contemporary cases reveal a discrepancy between the past aspirations to reconfigure work and the current realization, indicating that automation progresses faster in designers' discourses than in machines in construction robotics. The exploration pointed to the emergence of alternative approaches seeking robot support and worker–robot work-sharing for sensible integration of machines at ReAC and Canvas, next to bold visions of automation and productivity increase prevailing since the 1980s. Compared to the early aspirations of the civil engineers projecting deployments and worker displacement as a consequence at the DCE, the reconfiguration of work has so far emerged as the rearrangement of skills and a form of interdisciplinary labor (intertwining with a building trade) as the Canvas case displays.

7.4.1 Reflections on Technology Design in Construction Robotics

Additionally, the lines of investigation in this dissertation produced reflections on technology design in construction robotics, which I discuss in this section. Unlike in previous studies, the focus is not on providing direct implications for technically improving robotic systems (Xiao et al., 2022; Melenbrink et al., 2020) or addressing managerial barriers to their adoption in the construction industry (Chen et al., 2018; Delgado et al., 2019; Pan & Pan, 2020b; Park et al., 2023; Vora et al., 2024). Instead, I propose approaches and methods towards a sensible design and integration of these technologies into construction. To understand perspectives on robots and robot integration in the industry, previous studies have involved various stakeholders, including contractors, managers, engineers, and researchers (Delgado et al., 2019; Pan & Pan, 2020a, 2020b; Bademosi & Issa, 2021; Vora et al., 2024). A few exceptions aside (Kim et al., 2022; Park et al., 2023), these studies have not engaged with workers, who are potential end users and whose labor is likely to be affected directly.[12] Throughout this investigation, I listened to the perspectives of workers and technology designers and participated in or observed their work. The approaches and methods I offer essentially respond to four key issues I traced within these practitioners' perspectives and practices: the situated nature of tasks; skills, deskilling, and upskilling; the fitness of robots and sites; and roles for robots.

The Situated Nature of Tasks

An important point that some New Scaife Hall workers made concerns the gap between the current capabilities of robotic technologies and the complex, situated nature of their tasks. Consequently, these workers think that robots would invite additional challenges rather than have great utility in their tasks. These workers' focus on technology-job match and utility resonates with some other stakeholders' expectations for robots, reported as factors affecting robot adoption in previous studies. For instance, YeSeul Kim and coauthors (2022) showed that "perceived usefulness," "perceived relevance [of a technology to a job]," and "perceived output quality [by a technology]" are among the "contextual factors influencing" managers' and workers' perceptions of robot adoption in the United States.[13] In Kim and coauthors' (2022) study, these professionals see the utility of robots in "improv[ing] efficiency, productivity, or quality" while expressing negative views about these

12 While some of these studies, such as Pan and Pan (2020a) and Delgado and coauthors (2019), engaged contractors as participants, they do not specify whether those included workers. And while Vora and coauthors' (2024) study engaged superintendents, it does not specify whether it reached out to workers.

13 "Perceived usefulness" that Kim and coauthors (2022) used in this study is a variable that the researchers borrowed from the commonly referred "Technology Acceptance Model (TAM)" initially developed by Fred Davis (1989). In his article "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," Davis (1989) defines "perceived usefulness" as "the degree to which a person believes that using a particular system would enhance his or her job performance." "Perceived job relevance" and "perceived output quality" the researchers (2022) used are variables within the extended version of the model by Viswanath Venkatesh and Fred Davis (2000). In "A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies," Venkatesh and Davis (2000) define "perceived job relevance" as "individual's perception regarding the degree to which the target system is applicable to his or her job," and "perceived output quality" as the perception of "how well the system performs tasks [relevant to individual's job]."

technologies, deeming them to be “inadequate for dexterous and cognitive task[s]” or have “low adaptability to high variability.” Chintan Vora and coauthors (2024) reveal the influence of the same three factors on decision-making for robot adoption. According to their study, the “ability of the robot” to offer “improved safety, improved quality, or improved productivity” is a motivation for adopting robot use across a mix of stakeholders in the United States, including engineers, managers, and superintendents (Vora et al., 2024).

Based on the situated nature of their tasks, the New Scaife Hall workers do not really believe that robots are close to demonstrating this ability in practice. An overlooked task in furthering that ability is to better understand building tasks to better inform the design of robots and facilitate a more sensible integration of these technologies into sites. This requires stepping out of the lab, seeing sites, and talking to workers. Yet, how to gain a grounded understanding of construction work, the needs of workers, and the conditions on sites is an underrated line of inquiry compared to how to improve the models and machines in construction robotics, even though the first could immensely contribute to the latter. As I showed with the 1980s CMU, ReAC, and Canvas cases, the complex, layered makeup of construction work has been acknowledged since the field’s early years in the United States. ReAC and Canvas’s technology design processes are recent examples of close engagements with the work, workers, and sites towards what Mi Pan and Wei Pan (2020a) call “task-technology integrative and iterative design for construction robots.” Although these cases—and perhaps others—may indicate a beginning, it is difficult to say that the efforts to develop or experiment with methods for engaging with the work, workers, and sites in technology design are sufficient. Neither are the efforts to develop participatory design workflows in which technology designers, workers, and other industry stakeholders work together.

Therefore, next to advancing the models and machines, taking these technologies closer to delivering on workers’ and other stakeholders’ expectations is equally crucial. This may ultimately lead to more positive perceptions of robot adoption. To achieve this, in addition to the high-level understanding of tasks evident at the DCE in the 1980s and still prevalent in construction robotics today, situated engagements with construction work, workers, and sites, like in ReAC and Canvas, deserve more attention as an alternative technology design instrument. Further explorations of on-site research methods, in particular ethnography, may help establish these engagements and provide the technology design circle with a grounded understanding of construction work, the needs of workers, and the conditions on sites. Similarly, participatory design methods may help bring technology designers closer to construction and construction stakeholders in the design of technologies.

Skills, Deskilling, and Upskilling

Skills, deskilling, and upskilling are important issues in robot adoption, which I discussed in all cases, from conversations with the New Scaife Hall workers and fellow researchers in the ReAC project to observations of robotic finishing workflows at Canvas. Previous studies indicated that the perceived lack of technology skills and upskilling needed for robot use are barriers to robot adoption in stakeholders’ minds. According to Piyush Pradhananga and coauthors’ (2021) study, the perception of an “unskilled workforce” is considered to have a “moderate impact” on adoption by a group of construction professionals in the United States. In a similar study by Fopefoluwa Bademosi and Raja Issa (2021) also focusing on the United States, “low technology literacy” is a commonly mentioned barrier by construction technology professionals, managers, and researchers. In Chintan Vora and coauthors’ (2024) study, the “availability of skilled manpower to use the robotic technology” is identified as a factor in decision-making for robot adoption in the United States. Upskilling, which is

seen as a barrier from a managerial perspective in these studies, is a cause of "technological anxiety" on the workers' end, as YeSeul Kim and coauthors (2022) have shown. In their study, workers were more concerned that additional skill acquisition may be required "to adjust for automation" than about deskilling or a potential loss of value of their skills (Kim et al., 2022).[14]

In contrast, the New Scaife Hall workers who touched on the issue of skills in our conversations worried more about deskilling than upskilling for the near future. The potential of systems to lower the value of trade skills is an important point that emerged from these conversations. It is connected to these workers' comprehension of the parts of their tasks or skills that may or may not easily be performed with robotic systems. Conversely, finishers at Canvas did not worry about deskilling even though Basil facilitates entry into a particular step of drywall finishing because they recognized that their labor and tacit knowledge are indispensable to fully completing a job with the machine. Instead of anxiety, they displayed acceptance and competence in learning the new skills needed to work with the machine. These findings, combined with the managers' and workers' perspectives in previous studies, indicate that deskilling and upskilling may be perceived positively or negatively and, consequently, may work towards or against robot adoption.

Because technological disruption and labor displacement have been presented as unavoidable consequences since the 1980s, the field of construction robotics has mostly neglected to ask how human-centered technologies and workflows that do not discard but allow space for existing skills can be designed. How can the field position deskilling and upskilling as technology design considerations in developing such technologies? A potential first step could be to once again talk to workers to understand their perceptions of their own work, including the skilled parts in which they invest time to train for and find meaning in or which they leverage for economic reward, as well as those they regard as drudgery bringing no particular social or economic benefits. Ethnographic or participatory design research methods can be useful here. Based on this knowledge, the design of a particular system could be centered around the question of which steps or skills in a task make sense to perform with a machine and which ones should remain with workers, as I showed in the ReAC and Canvas cases. This question, which is often addressed from technical and economic perspectives, should also be attended to from the workers' perspective.

The Fitness of Robots and Sites

Despite optimism regarding the robotization of construction at the DCE in the 1980s, after forty years, the New Scaife Hall team voiced doubts as to whether robotic technologies are currently mature and reliable enough to work on sites and whether sites have the level of organization necessary for these technologies to operate. Similar doubts exist among other stakeholders. "Immature technology," as described by Piyush Pradhananga and coauthors (2021), is not only the "most-mentioned barrier" in construction robotics publications but is also rated by construction professionals as "having a moderate impact" on robot adoption in the United States. In line with the publications the researchers (2021) surveyed, "technology infancy and limitations" are mentioned as barriers to robot adoption frequently by construction technology professionals, managers, and researchers in Fopefoluwa Bademosi and Raja Issa's (2021) study. Some of these participants also referred to the "dynamic and unstructured nature of construction sites and processes" as an "operational" barrier in the United States (Bademosi & Issa, 2021). The workers involved in YeSeul Kim and coauthors' (2022) study stated that robots may hinder their work as they may cause "work-practice reconfigurations" and "congested project areas" on sites. Further, the managers interviewed

14 Still, the participants in Kim and coauthors' (2022) study who do carpentry or finishing tasks—such as plastering, drywall, or MEP—care about deskilling.

by the researchers (2022) considered the essential rearrangements on sites, including “new staff” or “necessary equipment” for deployment, operation, and maintenance, a challenge.

Time has shown that the adoption of robots has not been as easy as the engineers at the DCE hoped in the 1980s and that success is not solely contingent on technological advancements. Previous studies revealed that neither workers nor other industry stakeholders are necessarily enthusiastic about those advancements. Instead, they are pragmatically, and perhaps rightfully, interested in the result, that is, whether robots can operate and deliver the projected benefits in the real world. This is visible in how Canvas’s Basil is accepted in the industry.

Thus, although research to advance technologies is the primary line of work in construction robotics, the field should also incorporate site-based, contextual research on how to adapt these technologies to sites and sites to these technologies. Two approaches may be explored in this regard. The first is to leverage the existing models and machines in applications alternative to the automation of primary building tasks and in worker–robot scenarios different from dyadic worker–robot configurations. ReAC’s tool and hardware delivery robot, which takes over an ancillary task, is an example. One side of this approach is related to researching alternative opportunities for realizing these applications and scenarios, for which on-site research methods and ethnography may again be useful. The other side of this approach concerns shaping technologies according to site constraints, collecting data on sites to train models, and putting robots on sites to experimentally study their effects. The second approach consists in investing research efforts in making sites more suitable for robots while working to make robots more suitable for sites. This requires studies addressing how sites can be reconfigured for robots to operate with their current limitations. It may include studies on the scheduling of worker–robot sequences as well as the organization of work and material zones and the arrangement of markers or paths for specific technologies to work conveniently.

Roles for Robots

To increase productivity in the US construction industry, single-task robots with minimal to no human intervention were a hope chased after at the DCE in the 1980s. The engineers were aware of the long road ahead. Four decades later, some New Scaife Hall workers had little interest in robots due to the complications these technologies may bring into their tasks. Although they did not see robots executing tasks fully autonomously as a realistic near-term possibility, some were open to the idea of robots assisting with or taking over menial tasks such as cleaning. Other industry actors have both similar and different views on the role of robots. The majority of stakeholders who participated in Chintan Vora and coauthors’ (2024) study preferred some degree of autonomy with human involvement.[15] Interestingly, these participants supported the idea of one worker working with robots they consider easier to use, such as a rebar tying robot, and a team of workers working with robots they consider more difficult to use, such as a robotic excavator (Vora et al., 2024). The majority of workers who perform heavier concrete or steel tasks in YeSeul Kim and coauthors’ (2022) study did not believe that robots can handle their tasks and, therefore, found “assistive role[s]” for robots “more appropriate.”[16] If robots could perform full tasks, the majority of these participants would be open to “high autonomy” and shifting to “supervisory role[s]” (Kim et al., 2022). In contrast, half of the carpentry or finishing workers the researchers (2022) engaged, who carry out tasks such as

15 The participants in Vora and coauthors’ (2024) study include engineers, managers, superintendents, controllers, students, and virtual design and construction professionals. Close to eighty percent of these participants prefer some degree of robot autonomy with human involvement.

16 Kim and coauthors (2022) use “assistive role” to identify systems “designed to perform a subtask to help complete the main tasks,” whereas “task-executing role” to identify systems “designed to fully take over a main trade work that is directly related to the construction of a permanent building part.”

plastering, drywall, or MEP, valued a “task-executing role” for robots because they felt that robots could “improve productivity and quality” in their tasks. In addition, although these participants did not see any part of their tasks that a robot could assist with, the majority of them still favored “low autonomy” and an interactive workflow (Kim et al., 2022).

Fully autonomous, task-executing robots have been an ambitious proposition that has not turned into reality. Nor it is seen by workers as a near-term possibility. In this context, the specific role a particular robot could play in a building task, in conception or practice, and the kind of worker–robot scenario become important technology design considerations. For a more realistic integration of robots, tuning down expectations of full autonomy could be useful, instead focusing on finding opportunities and scenarios, or part of tasks, where current technologies might fit and be of use. As previously highlighted, much of ReAC’s and Canvas’s technology design processes revolved around this consideration. It was about finding a suitable role for a robot to perform. In the ReAC project, the role imagined and demonstrated with the tool and hardware delivery robot is one where robots become an instrument of support, a companion working as part of a worker team, rather than a substitute in the primary activities to which workers assign value.[17] A state-of-the-art reinforcement learning model and autonomous navigation are leveraged to provide support to workers. At Canvas, the role imagined and demonstrated with Basil is one where robots share part of a task with workers. Deep learning and computer vision are used to semi-autonomously perform one step in a task between the other steps performed by workers. Once again, understanding construction work, the conditions of sites, and the needs of workers through on-site research in construction robotics could offer a way to produce similar worker–robot scenarios and technologies.

7.4.2 Reflections on Pedagogy in Architecture, Construction, and Robotics

This dissertation is built on the premise that the development of robotic systems is a sociotechnical process contingent on technical knowledge and instruments, such as computer code, models, or hardware, together with the actors who participate and their conceptions and aspirations. By extension, these technologies are not just research instruments in the hands of these actors in labs or tech companies but become work artifacts in the hands of workers on sites and affect the way they shape buildings. Incorporating concepts and methods for understanding these premises into education is beneficial to raise future cohorts serving technological or intellectual production in architecture, construction, or robotics.

In architecture and construction education, this means providing students with a richer understanding of robotic technologies as objects of study, in addition to being a creative design, fabrication, or research instrument. In simple words, it means cultivating interest and excitement among those who would be interested in thinking, reading, and writing about robotic technologies and labor concerned with their design; next to hands-on working with these technologies. Within technology design education in HCI and robotics, this means guiding future designers to develop sensibilities based on the relationship between technologies and labor, and teaching them methods to engage laborers and other stakeholders in design. In interdisciplinary programs that combine both, a combination of these tactics may possibly be applied and experimented with. Putting these into practice could be through courses teaching concepts and methods of how to critically think about these technologies and how to develop and conduct research on them, next to courses teaching how to think and work with these technologies creatively or courses teaching technical research. These courses may enable students to produce rich, intriguing questions on technology

17 A similar proposition of using “immediate or near-future” technologies for “simple and small subtasks” is also offered by Kim and coauthors (2022).

design, technologies, and technologies' effects on work; and practice methods to form inquiry and investigate. They may enable them to develop thoughts or guardrails towards sensible, ethically informed technologies.

In a seminar course that I prepared and taught at the University of Pennsylvania's Weitzman School of Design, I practiced such pedagogy by transposing the inquiry produced for this dissertation into the context of MSD-RAS, a one-year post professional program. Broadly focusing on robotic fabrication, design computation, and material exploration, students in the program acquire skills and agency to "creatively engage" with technology and materials for building design through studios, technical courses, and seminars (MSD-RAS, 2019).[18] I was given an opportunity to redesign and teach the program's required RAS Theory seminar, aimed at providing a historical and conceptual foundation for the hands-on design and fabrication component. I put together a fresh version of the seminar, "Histories-Reconfigurations-Representations-Reflections," organized and titled based on the themes underlying the questions I explored in this dissertation. I taught this new version in the fall of 2023.

The new scope stemmed from the premise I mentioned earlier in this section, that is, CAD, AI, and robotic systems intersect with building design and construction practices and shape the way we design, represent, and build. Similar to the way I followed in the dissertation, it offers students an introduction to the intellectual and technical origins as well as the current configuration of this intersection from a sociotechnical angle. Through readings, writings, and conversations, students explored a series of key concepts in design, computation, and scholarly works in the seminar's four modules to comprehend and formulate new questions about this intersection (see Appendix E). In the first module, "Histories," we focused on the sites where AI/robotic systems were developed and converged with building design. Then, we reflected on the disposition and the effects of these systems in practice in "Reconfigurations" and on the materiality of these systems and the material fabrication practices they help facilitate in "Representations." The final module, "Reflections," served to discuss broader concepts, including machine learning and creativity, the politics of making, technological determinism, and AI ethics.

While the themes that give the modules their names are embedded in the questions in this dissertation, the students also explored concepts that are foundational to the descriptions and discussions in the chapters, such as sociotechnical effects, inscription, de-scription of technical objects, reconfiguration of practices, hand-eye-tool, skill, craft, and data representation (see Appendix E). Some of these concepts stimulated rich in-class discussions. For instance, the situated nature of building practices, which is widely described by scholars, mentioned by the New Scaife Hall workers, and discussed in the ReAC case was one of these. Similarly, the domain knowledge of building professionals, skills and deskilling, which are also discussed in the ReAC and Canvas cases, were reflected on. The question of technical formalization, representations of form, skill, material, and labor, and human–machine workflows, which are embedded in the CMU, ReAC, and Canvas studies, were also examined. Lastly, the role of machines and the representations shaping how we produce buildings, which underpin the Canvas study, were also part of the conversation.

Building on these weekly discussions, students were ultimately expected to formulate new questions to explore briefly in a final paper and presentation (see Appendix E). In light of the histories and key concepts of the convergence of AI and robotic systems and construction, a group of students worked on analyzing and reflecting on the recent body of work. Two surveyed the landscape of construction

18 See more on ras.design.upenn.edu.

robotics in the past decade, one focusing on the research developments at the Swiss Federal Institute of Technology (ETH) Institute of Technology in Architecture and the University of Stuttgart Institute for Computational Design and Construction, and the other concentrating on applications in the US AEC industry. The other student in this group focused on the concept of “craft” and explored the ways knowledge of craft “inform[s] contemporary design and digital fabrication of timber” in a few selected projects.

A couple of students zoomed in on the past. One investigated the ways early twentieth-century modular architecture relates to computation, and the other discussed whether off-the-shelf generative AI applications facilitate “new aesthetics in architecture” to reflect on these applications’ “true potentials and shortcomings” for the field. Another group ended up critically reflecting on their own work in the program through some of the key concepts explored. Drawing from AI ethics, one explored “accessibility and equity issues” in robotic fabrication education. Building on the concept of reconfiguration, another traced “the affordances, frictions, or limitations of working with code during collaborative making” based on a series of their own project collaborations in the program. The last one in this group discussed the importance and challenges of “prototyping customized fabrication tools and workflows” for designers.

Overall, my aim in the seminar was to cultivate agency and literacy in students so that they may comfortably navigate a cross-section of design, computing, and scholarly fields—as I did in this dissertation—and formulate questions similar to the ones I investigated. The quality of the students’ questions above, in my view, shows that this practice was successfully translated. While the dissertation guided the seminar, the seminar guided the students in producing those questions. This shows to me the generative potential of this dissertation.

7.4.3 Reflections on the Configuration of Ethnography

Two specific aspects related to the ethnographies in this dissertation deserve further attention. The first one concerns my dual role as an active participant and an observer within the ethnography of ReAC. As I described in Chapters 3 and 4, the object of study was the research group and its technology design process that I was part of, and the process included another ethnographic study (focusing on construction work) that I carried out.

This autoethnographic approach within the ReAC study is certainly distinct from being a neutral observer as I was at Canvas. The challenge was actively working as part of the group and simultaneously investigating the questions of this dissertation in our work. The value of the approach was that it allowed me to sensitize the group to particular aspects of construction. Through the ethnographic study, I had influence over the way I and my colleagues view and understand construction work, which then informed the way we conceived and developed a robotic system. The autoethnographic approach enabled me to reflect on this influence from a member’s view from within rather than that of an ethnographer looking in from outside.

The other important aspect is the incorporation of (technology) design and computing into the ReAC and Canvas fieldworks, expanding the capacities of conventional data collection, representation, and analysis. Each research tool that I described in Chapters 4 and 6 exemplifies a specific way of how design and computing can be leveraged to serve ethnography and produce ethnographic insights.

As the interface between a construction site and a research team, Prototype A, the data viewer, is an example of how the ethnographer’s design intervention can condition the exchange between the

work practice they study and the technology design practice that this study informs. More specifically, the practical contribution of this tool consisted in helping me consolidate a large quantity of dispersed and diverse ethnographic data sets on a single UI that can be navigated interactively. With this capability, the tool's contribution to the ReAC's system development process was that it exposed the team to the New Scaife Hall site and augmenting our collective discussions while conceiving a system with visual information. In short, it assisted me in "making the strange familiar" to myself and the rest of the group.[19] Alongside the field notes, the data organized via this tool helped me synthesize the task descriptions that I presented in Chapter 3.

Together with the software that Canvas uses to monitor their robots, Prototype B, the workflow visualization tool, helped extend the observational data collection and analysis into software space, where a major portion of robotic targeted spraying happens as a worker–machine feedback loop. The specific contribution of this tool lies in helping to organize this hidden loop, which comprises hundreds of workers' UI actions and machine responses, as computable data and giving it another visible form that is easily legible and compact to navigate. The prototype supported the second digital observation of the job site pilot test of Basil, in which I immersed myself in the UI sequence of an in-house finisher, following the in-person observation of the event. It offered a bird's-eye map of the loop enabling high-level exploration. This exploration produced the points I brought up about new, software-related skills around Basil, which I discussed in Chapter 6.

Lastly, Prototype C, the network graph visualization tool, is a design intervention that facilitated an analysis of work at Canvas. The ideation of this tool initiated a creative way of data collection from a collaborative work application. During my initial weeks at Canvas, the tool helped me get acquainted with the teams and remotely document their scope of work. It also helped me plan my involvement and improvise: by generating a map of the organization, I was able to choose its parts to zoom on in, participate, and observe each week. Further, the prototype expanded data representation and analysis. It enabled me to trace and visualize the interactions within and across the engineering sub-teams for ten weeks, along with each interaction's topic. This made the complex, interdisciplinary, and labor-intensive organization behind developing and deploying a machine for construction visible. In doing so, it helped me build the discussion in Chapter 6, in which I examined a range of distinct interactions and roles around Basil as an emerging technical labor in construction.

7.4.4 A Final Note

Technology design is not merely the practice of producing technical knowledge and artifacts enabling technologies. Producing knowledge about this practice that informs it back is a less explicit and perhaps underestimated endeavor. This is what I hoped to achieve with this dissertation. Looking into the state of construction robotics in past and present cases, I critically examined "construction work" in robotics and shed light on some of the sociotechnical changes occurring in construction work around robots. I provided a historically and socially grounded understanding of how we conceive robots and of the related reconfiguration of work, which, before this research, I imagined would unfold new considerations at the intersection of robotics and building practices. In this dissertation, I managed to produce some such considerations regarding technology design and pedagogy. With this, I hope that this dissertation will guide current and future technology designers, researchers, and scholars in producing new approaches, technologies, inquiries, or pedagogies. More importantly, I hope that it will sow new intriguing questions in every reader's mind.

19 The full version of this quote, "making the strange familiar and the familiar strange," is a common phrase anthropologists use to describe their field's mode of operating.

Bibliography

Bibliography

ABB Group (2021). ABB Robotics Advances Construction Industry Automation to Enable Safer and Sustainable Building. [online] Available at: https://new.abb.com/news/detail/78359/abb-robotics-advances-construction-industry-automation-to-enable-safer-and-sustainable-building [accessed 10 June 2024].

Abioye, S. O., Oyedele, L. O., Akanbi, L., Ajayi, A., Delgado, J. M. D., Bilal, M., Akinade, O. O., and Ahmed, A. (2021). Artificial Intelligence in the Construction Industry: A Review of Present Status, Opportunities and Future Challenges. *Journal of Building Engineering*, *44*, pp.103299.

Acemoglu, D. and Restrepo, P. (2022). Tasks, Automation, and the Rise in US Wage Inequality. *Econometrica*, *90*(5), pp.1973–2016.

Agin, G. (1972). *Description and Representation of Curved Objects*. (Doctoral dissertation, Stanford University).

Akrich, M. (1992). The De-scription of Technical Objects. In Bijker, W. E. and Law, J. (eds.), Shaping Technology / Building Society: Studies in Sociotechnical Change. MIT Press, pp.205–224.

Almono Limited Partnership (2022). The Site's History. [online] Available at: https://www.almono.org/thehistory [Accessed 10 November 2022].

Ames, M. (2019). *The Charisma Machine: The Life, Death, and Legacy of One Laptop Per Child*. MIT Press.

Andrychowicz, O. M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., and Schneider, J. (2020). Learning Dexterous In-Hand Manipulation. In *The International Journal of Robotics Research*, *39*(1), pp.3–20.

Apolinarska, A. A., Kuhn, M., Gramazio, F., and Kohler, M. (2021). Performance-Driven Design of a Reciprocal Frame Canopy: Timber Structure of the FutureTree. In *Proceedings of the 39th eCAADe Conference*, Volume 1, pp.497–504.

Argall, B. D., Chernova, S., Veloso, M., and Browning, B. (2009). A Survey of Robot Learning From Demonstration. In *Robotics and Autonomous Systems*, *57*(5), pp.469–483.

Armstrong, G., Gilge, C., Max, K., and Vora, S. (2023). KPMG 2023 Global Construction Survey: Familiar Challenges – New Approaches. KPMG International.

Associated General Contractors of America (2022). Construction Workforce Shortages Risk Undermining Infrastructure Projects as Most Contractors Struggle to Fill Open Positions. [online] Available at: https://www.agc.org/news/2022/08/31/construction-workforce-shortages-risk-undermining-infrastructure-projects-most-contractors-struggle [Accessed 6 June 2024].

Atanasoski, N. and Vora, K. (2019). *Surrogate Humanity: Race, Robots, and The Politics of Technological Futures*. Duke University Press.

Ayres, R. and Miller, S. (1983). *Robotics: Applications and Social Implications*. Ballinger.

Ayres, R. and Miller, S. (1981). *The Impact of Industrial Robots*. The Robotics Institute and Department of Engineering and Public Policy, Carnegie Mellon University.

Baecker, R. M. (2008). Timelines: Themes in the Early History of HCI—Some Unanswered Questions. *Interactions*, *15*(2), pp.22–27.

Ballard, L. A., Šabanović, S., Kaur, J., and Milojevic, S. (2012). George Charles Devol, Jr. [History]. In *IEEE Robotics & Automation Magazine*, *19*(3), pp.114–119.

Bard, J., Cupkova, D., Washburn, N., and Zeglin, G. (2018). Robotic Concrete Surface Finishing: A Moldless Approach to Creating Thermally Tuned Surface Geometry for Architectural Building Components Using Profile-3D-Printing. *Construction Robotics*, *2*(1-4), pp.53–65.

Bard, J., Blackwood, D., Sekhar, N., and Smith, B. (2016). Reality is Interface: Two Motion Capture Case Studies of Human–Machine Collaboration in High-Skill Domains. *International Journal of Architectural Computing*, *14*(4), pp.398–408.

Bard, J., Gannon, M., Jacobson-Weaver, Z., Jeffers, M., Smith, B., and Contreras, M. (2014). Seeing is Doing: Synthetic Tools for Robotically Augmented Fabrication in High-Skill Domains. In *Proceedings of the 34th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pp.409–416.

Barker, N., Jewitt, C., and Price, S. (2020). Becoming in Touch with Industrial Robots through Ethnography. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp.128–130.

Beane, M. I. (2020). In Storage, Yet on Display: An Empirical Investigation of Robots' Value as Social Signals. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp.83–91.

Beer, J. M., Fisk, A. D., and Rogers, W. A. (2014). Toward a Framework for Levels of Robot Autonomy in Human-Robot Interaction. In *Journal of Human-Robot Interaction*, *3*(2), pp.74–99.

Benanav, A. (2020). The Automation Discourse. In Benanav, A., *Automation and the Future of Work*. Verso Books, pp.1–13.

Bengfort, B., Bilbro, R., and Ojeda, T. (2018). *Applied Text Analysis with Python: Enabling Language-Aware Data Products with Machine Learning*. O'Reilly.

Bhat, P. and Zhao, Y. (2022). "I was Confused by It; It was Confused by Me:" Exploring the Experiences of People with Visual Impairments around Mobile Service Robots. *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), pp.1–26.

Bijker, W. E. and Law, J. (eds.) (1992). *Shaping Technology/Building Society: Studies in Sociotechnical Change*. 2nd printing, MIT Press, 1997.

Bijker, W. E., Hughes, T. P., and Pinch, T. J. (eds.) (1987). *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. 4th printing, MIT Press, 1993.

Billard, A. and Grollman, D. (2013). Robot Learning by Demonstration. [online] Available at: http://www.scholarpedia.org/article/Robot_learning_by_demonstration [Accessed 19 September 2021].

Bird, S., Klein, E., and Loper, E. (2009). *Natural Language Processing with Python*. O'Reilly.

Blaha, J. R., Lamoreux, J. P., and McKee, K. E. (1986). *Higher Order Languages for Robots*. Manufacturing Technology Information Analysis Center (MTIAC).

Blomberg, J. and Karasti, H. (2013). Reflections on 25 Years of Ethnography in CSCW. In *Computer Supported Cooperative Work (CSCW)*, 22(4-6), pp.373–423.

Blomberg, J., Suchman, L. and Trigg, R. H. (1996). Reflections on a Work-oriented Design Project. *Human–Computer Interaction*, *11*(3), pp.237–265.

Blond, L. (2019). Studying Robots Outside the Lab: HRI as Ethnography. *Paladyn, Journal of Behavioral Robotics*, *10*(1), 117–127.

Bloor, M. and Wood, F. (2006). Ethnography. In Bloor, M. and Wood, F., *Keywords in Qualitative Methods: A Vocabulary of Research Concepts*. Sage, pp.69–75.

Bademosi, F. and Issa, R. R. (2021). Factors Influencing Adoption and Integration of Construction Robotics and Automation Technology in the US. *Journal of Construction Engineering and Management*, *147*(8), p.04021075.

Bock, T. (2007). Construction Robotics. In *Autonomous Robots*, *22*(3), pp.201–209.

Bock, T. and Langenberg, S. (2014). Changing Building Sites: Industrialisation and Automation of the Building Process. In *Architectural Design*, *84*(3), pp.89–99.

Bock, T. and Linner, T. (2016). *Construction Robots: Elementary Technologies and Single-task Construction Robots*. Cambridge University Press.

Bodenhagen, L., Fischer, K., Winther, T. S., Langedijk, R. M., and Skjøth, M. M. (2019). Robot Use Cases for Real Needs: A Large-Scale Ethnographic Case Study. In *Paladyn Journal of Behavioral Robotics*, *10*(1), pp.193–206.

Bolles, R. and Paul, R. (1973). *The Use of Sensory Feedback in a Programmable Assembly System*. Stanford Artificial Intelligence Laboratory, Stanford University.

Bostock, M. and Observable, Inc. (2024). D3.js. [online] Available at: https://d3js.org/ [accessed 3 July 2022].

Bostock, M., Ogievetsky, V., and Heer, J. (2011). D3: Data-driven Documents. *IEEE Transactions on Visualization and Computer Graphics*, *17*(12), pp.2301–2309.

Bourne, D. A. and Fox, M. S. (1983). Autonomous Manufacturing: Automating the Job Shop. In *The Robotics Institute 1983 Annual Research Review*. Carnegie Mellon University, pp.7–19.

Bourne, D. A. and Fussell, P. (1982). *Designing Programming Languages for Manufacturing Cells*. The Robotics Institute, Carnegie Mellon University.

Bowker, G. C. and Star, S. L. (1999). Introduction: To Classify is Human. In *Sorting Things Out: Classification and Its Consequences*. MIT Press, pp.1–32.

Bødker, S., Ehn, P., Knudsen, J., Kyng, M. and Madsen, K. (1988). Computer Support for Cooperative Design. In *CSCW 88, Proceedings of the Conference on Computer-Supported Cooperative Work*, pp.377–394.

Braumann, J. and Brell-Cokcan, S. (2015). Adaptive Robot Control: New Parametric Workflows Directly From Design to KUKA Robots. In *Proceedings of the 33rd eCAADe Conference*, Volume 2, pp.243–250.

Braverman, H. (1974). Machinery. In Braverman, H., *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New Edition, Monthly Review Press, 1998, pp.127–162.

Braverman, H. (1974). A Final Note on Skill. In Braverman, H., *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. New Edition, Monthly Review Press, 1998, pp.294–310.

Bright, J. R. (1958). The Impact of Automation on the Work Force. In Bright, J. R., *Automation and Management*. Division of Research, Graduate School of Business Administration, Harvard University, pp.170–197.

Brooks, R. (2020). Roboticists Give Their Definition of “Robot” (Audio). [online] Available at: https://robots.ieee.org/learn/what-is-a-robot/ [Accessed 19 September 2021].

Brooks, R. (2018). A Brave, Creative, and Happy HRI. In *ACM Transactions on Human-Robot Interaction (THRI)*, *7*(1), pp.1:1–1:3.

Brooks, R. (1991). Intelligence Without Representation. In *Artificial intelligence*, *47*(1-3), pp.139–159.

Brinkmann, S. (2008). Interviewing. In Given, L. M. (ed.), *The Sage Encyclopedia of Qualitative Research Methods*. Sage, pp.470–472.

Buchanan, B. G. (2005). A (Very) Brief History of Artificial Intelligence. In *Ai Magazine*, *26*(4), pp.53–60.

Bullock, D. M. and Oppenheim, I. J. (1992). Object-oriented Programming in Robotics Research for Excavation. *Journal of Computing in Civil Engineering*, *6*(3), pp.370–385.

Button, G. and Sharrock, W. W. (2009). *Studies of Work and the Workplace in HCI: Concepts and Techniques* (Vol. 2). Morgan & Claypool Publishers.

Cardoso Llach, D. (2019). Tracing Design Ecologies: Collecting and Visualizing Ephemeral Data as a Method in Design and Technology Studies. In Vertesi, J. and Ribes, D. (eds.), *DigitalSTS: A Field Guide for Science & Technology Studies*. Princeton University Press, pp.451–471.

Cardoso Llach, D. (2015). *Builders of the Vision: Software and the Imagination of Design*. Routledge.

Cardoso Llach, D., Kaltman, E., Erdolu, E., and Furste, Z. (2021). An Archive of Interfaces: Exploring the Potential of Emulation for Software Research, Pedagogy, and Design. In *Proceedings of the ACM on Human-Computer Interaction (PACMHCI)*, *5*(CSCW2), pp.294:1–294:22.

Cardoso Llach, D., Bidgoli, A., and Darbari, S. (2017). Assisted Automation: Three Learning Experiences in Architectural Robotics. *International Journal of Architectural Computing*, *15*(1), pp.87–102.

Carnegie Mellon University (2024a). Robotics Program. [online] Available at: http://coursecatalog.web.cmu.edu/schools-colleges/schoolofcomputerscience/robotics/ [Accessed 10 July 2024].

Carnegie Mellon University (2024b). Cruise Controlled. [online] Available at: https://www.cmu.edu/about/index.html [Accessed 10 July 2024].

Carnegie Mellon University (2019). Mill 19 Opening Signals New Era of Manufacturing. [online] Available at: https://www.cmu.edu/news/stories/archives/2019/september/mill-19-grand-opening.html [Accessed 10 November 2022].

Carnegie Mellon University (2009). Remembering Three Mile Island. [online] Available at: https://www.cmu.edu/homepage/environment/2009/spring/30-years-ago.shtml [Accessed 10 September 2022].

Carnegie Mellon University College of Engineering (2022). Scaife Hall. [online] Available at: https://engineering.cmu.edu/about-us/giving/scaife-hall.html [Accessed 3 November 2022].

Carnegie Mellon University Robotics Institute (2022). History of the Robotics Institute. [online] Available at: https://www.ri.cmu.edu/about/ri-history/ [Accessed 10 September 2022].

Carnegie Mellon University Robotics Institute Youtube Channel (2019). RRV: Three Mile Island Rover Overview: 1984 CMU Robotics Institute Archive. [online] Available at: https://www.youtube.com/watch?v=bOaaA7tfSBw [Accessed 10 September 2022].

Carnegie Mellon University Robotics Institute Youtube Channel (2016). Terregator – CMU – Field Robotics Circa 1983. [online] Available at: https://www.youtube.com/watch?v=beGRM9ZmWlI [Accessed 10 September 2022].

Cha, E., Forlizzi, J., and Srinivasa, S. S. (2015). Robots in the Home: Qualitative and Quantitative Insights into Kitchen Organization. In *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.319–326.

Chang, W. L. and Šabanović, S. (2015a). Interaction Expands Function: Social Shaping of the Therapeutic Robot PARO in a Nursing Home. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. pp.343–350.

Chang, W. L. and Šabanović, S. (2015b). Studying Socially Assistive Robots in their Organizational Context: Studies with PARO in a Nursing Home. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, pp.227–228.

Chang, W. L. and Šabanovic, S. (2013). Potential Use of Robots in Taiwanese Nursing Homes. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.99–100.

Cheatle, A., Pelikan, H., Jung, M., and Jackson, S. (2019). Sensing (Co)operations: Articulation and Compensation in the Robotic Operating Room. In *Proceedings of the ACM on Human-Computer Interaction*, *3*(CSCW), pp.1–26.

Cheon, E., Schneiders, E., Diekjobst, K., and Skov, M. B. (2022a). Robots as a Place for Socializing: Influences of Collaborative Robots on Social Dynamics In-and Outside the Production Cells. In *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), pp.1–26.

Cheon, E., Schneiders, E., and Skov, M. B. (2022b). Working with Bounded Collaboration: A Qualitative Study on How Collaboration is Co-constructed Around Collaborative Robots in Industry. In *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), pp.1–34.

Cheon, E. and Su, N. M. (2018). Futuristic Autobiographies: Weaving Participant Narratives to Elicit Values Around Robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pp.388–397.

Cheon, E. and Su, N. M. (2016). Integrating Roboticist Values into a Value Sensitive Design Framework for Humanoid Robots. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.375–382.

Chen, Q., de Soto, B. G., and Adey, B. T. (2018). Construction Automation: Research Areas, Industry Concerns and Suggestions for Advancement. *Automation in Construction*, *94*, pp.22–38.

Chisholm, A. W. J. (1987). The Social Effects of Intelligent Manufacturing Systems. In *Robotics and Computer-Integrated Manufacturing*, *3*(2), pp.157–164.

Christensen, L. R. (2013). *Coordinative Practices in the Building Process: An Ethnographic Perspective*. Springer.

Christin, A. (2020). *Metrics at Work: Journalism and the Contested Meaning of Algorithms*. Princeton University Press.

Chun, B. and Knight, H. (2020). The Robot Makers: An Ethnography of Anthropomorphism at a Robotics Company. In *ACM Transactions on Human-Robot Interaction (HRI)*, *9*(3), pp.1–36.

Chun, W. H. K., Soon, W., Wardrip-Fruin, N., and Zhu, J. (2022). Software Studies, Revisited. A Roundtable on the Software Studies Series at MIT Press. *Computational Culture*, *9*.

Clearpath Robotics (2023). Husky Unmanned Ground Vehicle. [online] Available at: https://clearpathrobotics.com/husky-unmanned-ground-vehicle-robot/ [Accessed 28 December 2023].

Clearpath Robotics (2023). Jackal Unmanned Ground Vehicle. [online] Available at: https://clearpathrobotics.com/jackal-small-unmanned-ground-vehicle/ [Accessed 28 December 2023].

Clever, D., Hu, Y., and Mombaur, K. (2018). Humanoid Gait Generation in Complex Environments Based on Template Models and Optimality Principles Learned From Human Beings. In *The International Journal of Robotics Research*, *37*(10), pp.1184–1204.

Coghlan, S., Waycott, J., Lazar, A., and Barbosa Neves, B. (2021). Dignity, Autonomy, and Style of Company: Dimensions Older Adults Consider for Robot Companions. In *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW1), pp.1–25.

Colombino, T., Gallo, D., and Shreepriya, S. (2020). The Impact of Organizational Properties on the Design of Robotic Support for Employees with Disabilities. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp.166–168.

Congress of the United States (1982). *Japanese Productivity: Lessons for America —Hearing Before the Subcommittee on International Trade, Finance, and Security Economics of the Joint Economic Committee Congress of the United States*. U.S Government Printing Office.

Cousineau, L. and Miura, N. (1998). *Construction Robots: The Search for New Building Technology in Japan*. ASCE Publications.

Cutkosky, M. and Wright, P. K. (1986). Modeling Manufacturing Grips and Correlations with the Design of Robotic Hands. In *Proceedings of 1986 IEEE International Conference on Robotics and Automation*, Vol.3, pp.1533–1539.

Cyert, R. M. (1981). Prepared Statement of Richard M. Cyert. In *Japanese Productivity: Lessons for America —Hearing before the Subcommittee on International Trade, Finance, and Security Economics of the Joint Economic Committee Congress of the United States*. U.S. Government Printing Office, pp.56–64.

Cyert, R. M. (1978). The Management of Universities of Constant or Decreasing Size. In *Public Administration Review*, *38*(4), pp.344–349.

Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, pp.319–340.

de Soto, B. G. and Skibniewski, M. (2020). Future of Robotics and Automation in Construction. In Sawhney, A., Riley, M., and Irizarry, J. (eds.), *Construction 4.0*. Routledge, pp.289–306.

Delgado, J. M. D., Oyedele, L., Ajayi, A., Akanbi, L., Akinade, O., Bilal, M., and Owolabi, H. (2019). Robotics and Automated Systems in Construction: Understanding Industry-specific Challenges for Adoption. *Journal of Building Engineering*, *26*, p.100868.

Dautenhahn, K. (2018). Some Brief Thoughts on the Past and Future of Human-robot Interaction. In *ACM Transactions on Human-Robot Interaction (THRI)*, *7*(1), pp.4:1–4:3.

Deamer, P. (2012). BIM and Contemporary Labor. *Pidgin*, 15.

Dossick, C. S. and Neff, G. (2011). Messy Talk and Clean Technology: Communication, Problem-Solving and Collaboration Using Building Information Modelling. *The Engineering Project Organization Journal*, *1*(2), pp.83–93.

Dourish, P. and Bell, G. (2011). A Role for Ethnography: Methodology and Theory. In Dourish, P. and Bell, G., *Divining a Digital Future: Mess and Mythology in Ubiquitous Computing*. MIT Press, pp.61–89.

Dourish, P. and Button, G. (1998). On “Technomethodology": Foundational Relationships Between Ethnomethodology and System Design. *Human-computer Interaction*, *13*(4), pp.395–432.

Downs, A., Kootbally, Z., Harrison, W., Pilliptchak, P., Antonishek, B., Aksu, M., Schlenoff, C., and Gupta, S. K. (2021). Assessing Industrial Robot Agility Through International Competitions. In *Robotics and Computer-Integrated Manufacturing*, *70*, pp.102113.

Dusty Robotics (2023). The World’s Leading Robotic Layout System. [online] Available at: https://www.dustyrobotics.com/fieldprinter [Accessed 28 December 2023].

Duysburgh, P., Elprama, S. A., and Jacobs, A. (2014). Exploring the Social-technological Gap in Telesurgery: Collaboration within Distributed or Teams. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp.1537–1548.

Earnest, L. (1995). SAIL Away. [online] Available at: https://web.stanford.edu/~learnest/sail/sailaway.htm [Accessed 12 September 2022].

Earnest, L., McCarthy, J., Feigenbaum, E., and Lederberg, J. (eds.) (1973). *Final Report: The First Ten Years of Artificial Intelligence Research at Stanford*. Stanford Artificial Intelligence Laboratory, Stanford University.

Engelberger, J. F. (1980). *Robotics in Practice: Management and Applications of Industrial Robots*. Kogan Page & Avebury Publishing Company.

Escobar-Planas, M., Charisi, V., and Gomez, E. (2022). "That Robot Played with Us!" Children's Perceptions of a Robot after a Child-Robot Group Interaction. In *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), pp.1–23.

Feldman, G. (1967). Butterfinger (AI Film #42). [online] Available at: https://www.csail.mit.edu/node/6649 [Accessed 5 September 2022].

Filippi, G. and Theureau, J. (1993). Analyzing Cooperative Work in an Urban Traffic Control Room for the Design of a Coordination Support System. In *Proceedings of the Third European Conference on Computer-Supported Cooperative Work 13–17 September 1993, Milan, Italy ECSCW’93*, pp. 171–186.

Firth, C., Dunn, K., and Haeusler, M. H. (2020). Design Process for a Soft Flexible Palm: Improving Grasp Strength in an Anthropomorphic End Effector for Collaborative Robots in Construction. In *Proceedings of the 38th eCAADe Conference*, Volume 2, pp.423–432.

Fischer, F., Fischer, K., and Palinko, O. (2023). A Persuasive Hand Sanitizer Robot in the Wild: the Effect of Persuasive Speech on the Use of a Hand Sanitizer Robot. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pp.649–652.

Fisher, E. L. and Nof, S. Y. (1982). *The Work Abilities of Industrial Robots*. The School of Industrial Engineering, Purdue University.

Fleck, J. (1982). Development and Establishment in Artificial Intelligence. In Elias, N., Martins, H., and Whitley, R. P. (eds.), *Scientific Establishments and Hierarchies*. D. Reidel Publishing Company, pp.169-217.

Forlizzi, J. and DiSalvo, C. (2006). Service Robots in the Domestic Environment: A Study of the Roomba Vacuum in the Home. In *Proceedings of the 1st ACM SIGCHI / SIGART Conference on Human-Robot Interaction (HRI)*, pp.258–265.

Forsythe, D. (2001). *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford University Press.

Fuller, M. (2008). Introduction. In Fuller, M. (ed.), *Software Studies: A Lexicon*. MIT Press, pp.1–14.

Garcia del Castillo y López, J. L. (2019). Robot Ex Machina. In *Proceedings of the 39th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pp.40–49.

Gasparetto, A. and Scalera, L. (2019). A Brief History of Industrial Robotics in the 20th Century. In *Advances in Historical Studies*, *8*(1), pp.24–35.

Goodrich, M. A. and Schultz, A. C. (2007). *Human-Robot Interaction: A Survey*. Now Publishers Inc.

Google-Oxford Dictionary (2024). Work. [online] Available at: https://www.google.com/search?q=dictionary&rlz=1C1CHBF_enUS813US813&oq=dictionary&gs_lcrp=EgZjaHJvbWUqBggAEEUYOzIGCAAQRRg7MgYIARBFGDwyBwgCEAAYjwIyBwgDEAAYjwIyBggEEEUYOzIGCAUQRRg8MgYIBhBFGDwyBggHEEUYPNIBCDE4NTIqMGo0qAIAsAIB&sourceid=chrome&ie=UTF-8#dobs=work [Accessed 20 September 2021].

Goolsbee, A. and Syverson, C. (2023). *The Strange and Awful Path of Productivity in the U.S. Construction Sector* (No. w30845). National Bureau of Economic Research.

Grint, K. (2005). What is Work?. In Grint, K., *The Sociology of Work: Introduction*. Polity, pp.6–44.

Grudin, J. (2017). *From Tool to Partner: The Evolution of Human-Computer Interaction*. Morgan and Claypool.

Grudin, J., (1994). Computer-Supported Cooperative Work: History and Focus. In *Computer*, *27*(5), pp.19–26.

Grudin, J. (1990). The Computer Reaches Out: The Historical Continuity of Interface Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 261–268.

Gypsum Association (2021). GA-214-2021 – Levels of Finish for Gypsum Panel Products. [online] Available at: https://www.gypsumpublications.com/product/2213/levels-of-finish-for-gypsum-panel-products-pdf-download-ga-214-2021-pdf.

Gypsum Association (2015). GA-214-2015 – Recommended Levels of Finish for Gypsum Board, Glass Mat, and Fiber-Reinforced Gypsum Panels.

Haas, C. (1990). *A Model of Pavement Surfaces*. (Doctoral dissertation, Carnegie Mellon University).

Haas, C., Skibniewski, M., and Budny, E. (1995). Robotics in Civil Engineering. In *Computer-Aided Civil and Infrastructure Engineering*, *10*(5), pp.371–381.

Haas, C., Hendrickson, C., McNeil, S. and Bullock, D. (1992). A Field Prototype of a Robotic Pavement Crack Sealing System. In *Proceedings of the 9th International Symposium on Automation and Robotics in Construction (ISARC'92),* pp.313–322.

Hammersley, M. and Atkinson, P. (1983). *Ethnography: Principles in Practice*. 3rd edition, Routledge, 2007.

Harper, R. (2000). Analysing Work Practice and the Potential Role of New Technology at the International Monetary Fund: Some Remarks on the Role of Ethnomethodology. In Luff, P., Hindmarsh, J., and Heath, C. (eds.), *Workplace Studies: Recovering Work Practice and Informing System Design.* Cambridge University Press, pp.169–186.

Harris, K. (2020). Forty Years of Falling Manufacturing Employment. In *Beyond the Numbers: Employment & Unemployment 9*(16). U.S. Bureau of Labor Statistics.

Hasegawa, Y. (1983). Robotization of Construction Work. Robot No.38, Kikai Shimko, Tokyo, Japan.

Hasse, C. (2019). The Multi-variation Approach: Cross-case Analysis of Ethnographic Fieldwork. In *Paladyn, Journal of Behavioral Robotics*, *10*(1), pp.219–227.

Hatch, P., Rahman, M. A., and Michaelis, J. E. (2023). PATHWiSE: An Authoring Tool to Support Teachers to Create Robot-Supported Social Learning Experiences During Homework. *Proceedings of the ACM on Human-Computer Interaction*, *7*(CSCW1), pp.1–23.

Heath, C. and Luff, P. (2000). *Technology in Action*. Cambridge University Press.

Hendrickson, C. (2021). Interviewed by Emek Erdolu [Zoom]. 3 February.

Hilti Corporation (2020). Hilti Unveils Bim-enabled Construction Jobsite Robot. [online] Available at: https://www.hilti.group/content/hilti/CP/XX/en/company/media-relations/media-releases/Jaibot.html [Accessed 28 December 2023].

Holzer, D. (2009). *Sense-making Across Collaborating Disciplines in the Early Stages of Architectural Design*. (Doctoral dissertation, Royal Melbourne Institute of Technology).

Horst, R. L., Law, T. M., and Silverman, E. B. (1984). A Procedure for Developing a Standardized Robotic Maintenance System for Nuclear Power Plants. In *Robotics and Remote Handling in Hostile Environments*. American Nuclear Society.

Hsu, H. (2021). Oral History of Raj Reddy – Computer History Museum. 14 April. [online] Available at: https://archive.computerhistory.org/resources/access/text/2021/06/102792182-05-01-acc.pdf [Accessed 17 September 2022].

Hughes, J., King, V., Rodden, T., and Andersen, H. (1994). Moving Out From the Control Room: Ethnography in System Design. In *Proceedings of the 1994 ACM conference on Computer Supported Cooperative Work*, pp.429–439.

Huhtamo, E. and Parikka, J. (eds.) (2011). *Media Archaeology: Approaches, Applications, and Implications*. University of California Press.

Investopedia – The Investopedia Team (2024). US Recessions Throughout History: Causes and Effects. [online] Available at: https://www.investopedia.com/articles/economics/08/past-recessions.asp [Accessed 17 June 2024].

Investopedia – Young, J. (2023). Too Big to Fail: Definition, History, and Reforms. [online] Available at: https://www.investopedia.com/terms/t/too-big-to-fail.asp [Accessed 17 June 2024].

Jenny, S. E., Lloret-Fritschi, E., Gramazio, F., and Kohler, M. (2020). Crafting Plaster Through Continuous Mobile Robotic Fabrication On-Site. *Construction Robotics*, *4*(3), pp.261–271.

JHU Applied Physics Laboratory Youtube Channel (2012). APL Flashback: 1965 – The "Hopkins Beast" Mobile Automaton. [online] Available at: https://www.youtube.com/watch?v=zMpEeag7kkM [Accessed 28 September 2022].

Jocher, G., Nishimura, K., Mineeva, T., and Vilariño, R. (2020). YOLOv5 (2020). [online] Available at: https://github.com/ultralytics/yolov5 [Accessed 28 December 2023].

Johnson, L. B. (1964). Remarks Upon Signing Bill Creating the National Commission on Technology, Automation, and Economic Progress. [online by Gerhard Peters and John Woolley, the American Presidency Project] Available at: https://www.presidency.ucsb.edu/node/241917 [Accessed 17 June 2023].

Juel, W. K., Haarslev, F., Ramirez, E. R., Marchetti, E., Fischer, K., Shaikh, D., Manoonpong, P., Hauch, C., Bodenhagen, L., and Krüger, N. (2020). Smooth Robot: Design for a Novel Modular Welfare Robot. *Journal of Intelligent & Robotic Systems*, *98*(1), pp.19–37.

Jung, M. and Hinds, P. (2018). Robots in the Wild: A Time for More Robust Theories of Human-Robot Interaction. In *ACM Transactions on Human-Robot Interaction (HRI)*, *7*(1), pp.2:1–2:5.

Kain, P. (1998). How to Do a Close Reading. Writing Center at Harvard University. [online] Available at: https://writingcenter.fas.harvard.edu/pages/how-do-close-reading [Accessed 12 September 2021].

Kanade, T. (1987). Introduction. In *The Robotics Institute 1987 Annual Research Review*. The Robotics Institute, Carnegie Mellon University.

Khovanskaya, V., Sengers, P., Mazmanian, M. and Darrah, C. (2017). Reworking the Gaps Between Design and Ethnography. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp.5373–5385.

Kiesler, S. and Goodrich, M. A. (2018). The Science of Human-robot Interaction. In *ACM Transactions on Human-Robot Interaction (HRI)*, 7(1), pp.9:1–9:3.

Kim, H., Lee, H., Chung, S., and Kim, C. (2007). User-Centered Approach to Path Planning of Cleaning Robots: Analyzing User's Cleaning Behavior. In *Proceedings of the 2007 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.373–380.

Kim, J., Chung, D., Kim, Y., and Kim, H. (2022). Deep Learning-based 3D Reconstruction of Scaffolds Using a Robot Dog. In *Automation in Construction*, *134*, pp.104092.

Kim, Y., Kim, H., Murphy, R., Lee, S., and Ahn, C. R. (2022). Delegation or Collaboration: Understanding Different Construction Stakeholders' Perceptions of Robotization. *Journal of Management in Engineering*, *38*(1), pp.04021084.

King, M. T. (2012). Working With/In the Archives. In Gunn, S. and Faire, L. (eds.), *Research Methods for History*. Edinburgh University Press, pp.13–29.

Knorr-Cetina, K. (1983a). New Developments in Science Studies: The Ethnographic Challenge. *Canadian Journal of Sociology/Cahiers Canadiens de Sociologie*, pp.153–177.

Knorr-Cetina, K. (1983b). The Ethnographic Study of Scientific Work: Towards a Constructivist Interpretation of Science. In Knorr-Cetina K. and M. Mulkay M. (eds.), *Science Observed: Perspectives on the Social Study of Science*. Sage, pp.115–140.

Komiya, R. and Irino, M. (1992). Japan's Economic and Industrial Policy in the 1980's. In *Keynes and the Economic Policies of the 1980s*, pp.137–174.

Kortenkamp, D., Simmons, R., and Brugali, D. (2016). Robotic Systems Architectures and Programming. In Siciliano, B. and Khatib, O. (eds.), *Springer Handbook of Robotics*. Springer, pp.283–305.

Kraft, K. (2016). Robots Against Infectious Diseases. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.627–628.

Langdon, F. (1983). Japan–United States Trade Friction: The Reciprocity Issue. In *Asian Survey*, pp.653–666.

Larson, R. and Csikszentmihalyi, M. (2014). The Experience Sampling Method. In Csikszentmihalyi, M. (ed.), *Flow and the Foundations of Positive Psychology.* Springer, pp.21–34.

Latour, B. and Woolgar, S. (1979). *Laboratory Life: The Construction of Scientific Facts*. 2nd edition, Princeton University Press, 1986.

Lee, H. R., Cheon, E., De Graaf, M., Alves-Oliveira, P., Zaga, C., and Young, J. (2019). Robots for Social Good: Exploring Critical Design for HRI. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.681–682.

Lee, H. R. and Sabanović, S. (2014). Culturally Variable Preferences for Robot Design and Use in South Korea, Turkey, and the United States. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction*, pp.17–24.

Leite, I., Castellano, G., Pereira, A., Martinho, C., and Paiva, A. (2012). Modelling Empathic Behaviour in a Robotic Game Companion for Children: An Ethnographic Study in Real-world Settings. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, pp.367–374.

Liang, C. J., Kamat, V. R., and Menassa, C. C. (2020). Teaching Robots to Perform Quasi-Repetitive Construction Tasks Through Human Demonstration. In *Automation in Construction*, *120*, pp.103370.

Liu, K., Xiong, C. H., He, L., Chen, W. B., and Huang, X. L. (2018). Postural Synergy Based Design of Exoskeleton Robot Replicating Human Arm Reaching Movements. In *Robotics and Autonomous Systems*, *99*, pp.84–96.

Loukissas, Y. A. (2012). *Co-designers: Cultures of Computer Simulation in Architecture*. Routledge.

Macdonald, K. M. (1995). Sociological Analysis of the Professions. In Macdonald, K. M., *The Sociology of the Professions*. SAGE Publications, pp.1–35.

Maibaum, A., Bischof, A., Hergesell, J., and Lipp, B. (2021). A Critique of Robotics in Health Care. *AI & Society*, pp.1–11.

Marino, M. C. (2020). *Critical Code Studies*. MIT Press.

Markoff, J. (2018). A. M. Turing Award Oral History Interview with Raj Reddy. 26 February. [online] Available at: https://amturing.acm.org/pdf/ReddyTuringTranscript.pdf [Accessed 17 September 2022].

Marsh, A. (2022). In 1961, the First Robot Arm Punched In. [online] Available at: https://spectrum.ieee.org/unimation-robot [Accessed 17 June 2023].

Marsh, A. (2004). Tracking the PUMA. In *Proceedings of the IEEE Conference on the History of Electronics*.

Mast, M. (2014). *Human-Robot Interaction for Semi-Autonomous Assistive Robots: Empirical Studies and an Interaction Concept for Supporting Elderly People at Home.* (Doctoral dissertation, Linköping University).

Matarić, M. (2018). On Relevance: Balancing Theory and Practice in HRI. In *ACM Transactions on Human-Robot Interaction (THRI)*, *7*(1), pp.8:1–8:2.

Mavrogiannis, C., Baldini, F., Wang, A., Zhao, D., Trautman, P., Steinfeld, A., and Oh, J. (2023). Core Challenges of Social Robot Navigation: A Survey. *ACM Transactions on Human-Robot Interaction*, *12*(3), pp.1–39.

McCullough, M. (1996). Human Context. In McCullough, M., *Abstracting Craft: The Practiced Digital Hand*. MIT Press, pp.1–82.

McKechnie, L. E. F. (2008). Participant Observation. In Given, L. M. (ed.), *The Sage Encyclopedia of Qualitative Research Methods*. Sage, pp.598–599.

McKinney, W. (2013). *Python for Data Analysis*. O'Reilly.

McKinney, W. (2010). Data Structures for Statistical Computing in Python. In *SciPy* (Vol.445, No.1), pp.51–56.

Mehrbod, S., Staub-French, S., Mahyar, N., and Tory, M. (2019). Characterizing Interactions with BIM Tools and Artifacts in Building Design Coordination Meetings. In *Automation in Construction*, *98*, pp.195–213.

Merriam, S. B. (2009). Types of Qualitative Research. In Merriam, S. B., *Qualitative Research: A Guide to Design and Implementation*. Jossey-Bass, pp.21–38.

Meissner, A., Trübswetter, A., Conti-Kufner, A. S., and Schmidtler, J. (2020). Friend or Foe? Understanding Assembly Workers' Acceptance of Human-Robot Collaboration. *ACM Transactions on Human-Robot Interaction (THRI)*, *10*(1), pp.1–30.

Melenbrink, N., Kassabian, P., Menges, A., and Werfel, J. (2017). Towards Force-Aware Robot Collectives for On-Site Construction. In *Proceedings of the 37th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pp.382–391.

Melenbrink, N., Werfel, J., and Menges, A. (2020). On-Site Autonomous Construction Robots: Towards Unsupervised Building. In *Automation in Construction*, *119*, pp.103312.

Merriam-Webster Dictionary (2024). De-skill. [online] Available at: https://www.merriam-webster.com/dictionary/de-skill#:~:text=transitive%20verb,a%20job%20by%20(a%20worker) [Accessed 20 June 2024].

Merriam-Webster Dictionary (2024). Labor. [online] Available at: https://www.merriam-webster.com/dictionary/labor [Accessed 20 September 2021].

Miller, R. K. (1989a). Robots in Industry. In Miller, R. K., *Industrial Robot Handbook: Case Histories of Effective Robot Use in 70 Industries*. Springer, pp.1–9.

Miller, R. K. (1989b). The Robot Manufacturing Cell. In Miller, R. K., *Industrial Robot Handbook: Case Histories of Effective Robot Use in 70 Industries*. Springer, pp.61–69.

Miller, R. K. (1989c). Construction. In Miller, R. K., *Industrial Robot Handbook: Case Histories of Effective Robot Use in 70 Industries*. Springer, pp.673–678.

Miller, S. (2021). Interviewed by Emek Erdolu [Zoom]. 28 January.

Minsky, M. (1972). Mini-Robot Proposal to ARPA. Massachusetts Institute of Technology AI Laboratory.

Mitchell, W. (2017). Automation and Full Employment — Back to the 1960s. [online] Available at: https://billmitchell.org/blog/?p=37361 [Accessed 17 June 2023].

Moorman, A., Liu, J., and Sabin, J. E. (2016). RoboSense: Context-Dependent Robotic Design Protocols and Tools. In *Proceedings of the 36th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pp.174–183.

Moravec, H. P. (1983). The Stanford Cart and the CMU Rover. In *Proceedings of the IEEE*, *71*(7), pp.872–884.

Moravec, H. (1980). *Obstacle Avoidance and Navigation in the Real World by a Seeing Robot Rover*. (Doctoral dissertation, Stanford University).

Mossfeldt Nickelsen, N. C. (2019). Imagining and Tinkering with Assistive Robotics in Care for the Disabled. *Paladyn, Journal of Behavioral Robotics*, *10*(1), pp.128–139.

MSD-RAS (2019). Program Overview. [online] Available at: https://ras.design.upenn.edu/ [accessed 10 June 2024].

Mubin, O., Wadibhasme, K., Jordan, P. and Obaid, M. (2019). Reflecting on the Presence of Science Fiction Robots in Computing Literature. *ACM Transactions on Human-Robot Interaction (THRI)*, *8*(1), pp.1-25.

Murphy, K. M. and Marcus, G. E. (2013). Epilogue: Ethnography and Design, Ethnography in Design . . . Ethnography by Design. In Gunn, W., Otto, T., Smith, R. C. (eds.), *Design Anthropology: Theory and Practice*. Bloomsbury, pp.251–268.

Myers, B. A. (1998). A Brief History of Human-computer Interaction Technology. In *Interactions*, *5*(2), pp.44–54.

Nielsen, S., Skov, M. B., Hansen, K. D., and Kaszowska, A. (2023). Using User-Generated YouTube Videos to Understand Unguided Interactions with Robots in Public Places. In *ACM Transactions on Human-Robot Interaction (THRI)*, *12*(1), pp.1–40.

Nilsson, N. J. (1984). Shakey the Robot. Artificial Intelligence Center Computer Science and Technology Division. SRI International.

NLTK Team (2024). Natural Language Toolkit (NLTK). [online] Available at: https://www.nltk.org/index.html [accessed 3 July 2022].

Ochoa-Franco, L. A., Haas, C. T., Dailey, C. M., and Traver, A. E. (1994). Construction Automation Research Database (CARD). In *Automation and Robotics in Construction XI,* pp.523–530.

Oppenheim, I. (2021). Interviewed by Emek Erdolu [Zoom]. 26 January.

Oppenheim, I. and Skibniewski, M. (1988). Construction, Robots In. In Dorf, R. C. and Nof, S. Y. (eds.), *International Encyclopedia of Robotics: Applications and Automation, Volume 1*. John Wiley & Sons, pp.240–249.

Ostrowski, A. K., Breazeal, C., and Park, H. W. (2022). Mixed-Method Long-Term Robot Usage: Older Adults' Lived Experience of Social Robots. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.33–42.

Paluch, R. and Müller, C. (2022). "That's Something for Children" An Ethnographic Study of Attitudes and Practices of Care Attendants and Nursing Home Residents Towards Robotic Pets. *Proceedings of the ACM on Human-Computer Interaction*, *6*(GROUP), pp.1–35.

Pan, M. and Pan, W. (2020a). Stakeholder Perceptions of the Future Application of Construction Robots for Buildings in a Dialectical System Framework. *Journal of Management in Engineering*, *36*(6), p.04020080.

Pan, M. and Pan, W. (2020b). Understanding the Determinants of Construction Robot Adoption: Perspective of Building Contractors. *Journal of Construction Engineering and Management*, *146*(5), p.04020040.

Pan, Y. and Zhang, L. (2021). Roles of Artificial Intelligence in Construction Engineering and Management: A Critical Review and Future Trends. In *Automation in Construction*, *122*, pp.103517.

Pantofaru, C., Takayama, L., Foote, T., and Soto, B. (2012). Exploring the Role of Robots in Home Organization. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, pp.327–334.

Park, S., Yu, H., Menassa, C. C., and Kamat, V. R. (2023). A Comprehensive Evaluation of Factors Influencing Acceptance of Robotic Assistants in Field Construction Work. *Journal of Management in Engineering*, *39*(3), p.04023010.

Pelikan, H. R., Broth, M., and Keevallik, L. (2020). "Are You Sad, Cozmo?" How Humans Make Sense of a Home Robot's Emotion Displays. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp.461–470.

Pelikan, H. R., Cheatle, A., Jung, M. F., and Jackson, S. J. (2018). Operating at a Distance—How a Teleoperated Surgical Robot Reconfigures Teamwork in the Operating Room. In *Proceedings of the ACM on Human-Computer Interaction (PACMHCI)*, *2*(CSCW), pp.1–28.

Petrosky, L. J. and Oppenheim, I. J. (1989). Model Based Manipulator Control. In *JPL, California Institute of Technology, Proceedings of the NASA Conference on Space Telerobotics, Volume 4*.

Petrosky, L. J. and Oppenheim, I. J. (1988). Application of Model Based Control to Robotic Manipulators. In *NASA. Lyndon B. Johnson Space Center, 2nd Annual Workshop on Space Operations Automation and Robotics (SOAR 1988)*.

Pinelle, D. and Gutwin, C. (2000). A Review of Groupware Evaluations. In *Proceedings of IEEE 9th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE 2000)*, pp.86–91.

Plunkert, L. M. (1990). The 1980s: A Decade of Job Growth and Industry Shifts. In *Monthly Labor Review.* Bureau of Labor Statistics, pp.3–15.

Polanyi, M. (1966a). Tacit Knowing. In Stehr, N. and Grundmann, R. (eds.) (2005), *Knowledge: Critical Concepts (Volume II Knowledge and Society: Forms of Knowledge)*. Routledge, pp.101–114.

Polanyi, M. (1966b). *The Tacit Dimension*. The University of Chicago Press.

Polanyi, M. (1958). *Personal Knowledge: Towards a Post-Critical Philosophy*. Routledge.

Pradhananga, P., ElZomor, M., and Santi Kasabdji, G. (2021). Identifying the Challenges to Adopting Robotics in the US Construction Industry. *Journal of Construction Engineering and Management*, *147*(5), p.05021003.

Pratt, G. (2020). Roboticists Give Their Definition of "Robot" (Audio). [online] Available at: https://robots.ieee.org/learn/what-is-a-robot/ [Accessed 19 September 2021].

Processing – Fry B. and Reas C. (2001). Processing. [online] Available at: https://processing.org/ [Accessed 10 November 2022].

Randall, D., Harper, R., and Rouncefield, M. (2005). Fieldwork and Ethnography: A Perspective From CSCW. In *Ethnographic Praxis in Industry Conference Proceedings* (Vol.2005, No.1), pp.81–99).

Reddy, R. (1989). Looking to The Future. In *The Robotics Institute 1989 Annual Research Review*. The Robotics Institute, Carnegie Mellon University.

Rhee, J. (2018). *The Robotic Imaginary: The Human and The Price of Dehumanized Labor*. University of Minnesota Press.

Ritchie, D. A. (1994). *Doing Oral History: A Practical Guide*. Oxford University Press, 2003.

Rossi, G. and Nicholas, P. (2019). Haptic Learning: Towards Neural-Network-Based Adaptive Cobot Path-Planning for Unstructured Spaces. In *Proceedings of the 37th eCAADe and 23rd SIGraDi Conference*, Volume 2, pp.201–210.

Roth, B. (2016). Foreword. In Siciliano, B. and Khatib, O. (eds.), *Springer Handbook of Robotics*. Springer pp.V–IX.

Roth, B., Rastegar, J., and Scheinman, V. (1974). On the Design of Computer Controlled Manipulators. In First CISM — IFToMM Symposium - On Theory and Practice of Robots and Manipulators. Springer-Verlag Wien, pp.93–113.

Ryle, G. (1951). Thinking and Language. In G. Ryle, *Collected Papers. Volume II: Collected Essays, 1929–1968*. Hutchinson & Co, 1971, pp.258–271.

Šabanović, S. (2007). *Imagine All the Robots: Developing a Critical Practice of Cultural and Disciplinary Traversals in Social Robotics.* (Doctoral dissertation, Rensselaer Polytechnic Institute).

Sabelli, A. M., Kanda, T., and Hagita, N. (2011). A Conversational Robot in an Elderly Care Center: an Ethnographic Study. In *Proceedings of the 6th International Conference on Human-robot Interaction*, pp.37–44.

Saidi, K. S., Bock, T., and Georgoulas, C. (2016). Robotics in Construction. In Siciliano, B. and Khatib, O. (eds.), *Springer Handbook of Robotics*. Springer, pp.1493–1520.

Saldana, J. (2013). *The Coding Manual for Qualitative Researchers*. Sage.

Santos, P. M. M., Campilho, R. D. S. G., and Silva, F. J. G. (2021). A New Concept of Full-Automated Equipment for the Manufacture of Shirt Collars and Cuffs. In *Robotics and Computer-Integrated Manufacturing*, *67*, pp.102023.

Scheinman, V. (2011). Stanford Hydraulic Arm (Artifact Details) [online] Available at: https://www.computerhistory.org/collections/catalog/102689927 [Accessed 9 September 2022].

Schmidt, K. (2011). The Concept of 'Work' in CSCW. In *Computer Supported Cooperative Work (CSCW)*, *20*(4), pp.341–401.

Schmidt, K. and Bannon, L. J. (2013). Constructing CSCW: The First Quarter Century. In *Computer Supported Cooperative Work (CSCW)*, *22*(4-6), pp.345–372.

Schmidt, K. and Wagner, I. (2004). Ordering Systems: Coordinative Practices and Artifacts in Architectural Design and Planning. *Computer Supported Cooperative Work (CSCW)*, *13*(5), pp.349–408.

Schmidt, K. and Wagner, I. (2003). Ordering Systems: Coordinative Practices in Architectural Design and Planning. In *Proceedings of the 2003 ACM International Conference on Supporting Group Work*, pp.274–283.

Schmidt, K. and Wagner, I. (2002). Coordinative Artifacts in Architectural Practice. In *Proceedings of the Fifth International Conference on the Design of Cooperative Systems (COOP)*, Vol. 2, pp.257–274.

Shim, S. J. and Kumar, A. (2005). Selection and Diversification of Market Segments for Robotics Products. *Industrial Robot: An International Journal*. Emerald Group Publishing, pp.401–407.

Shin, K. G. and Bonner, S. (1982). *A Comparative Study of Robot Languages*. Center for Robotics and Integrated Manufacturing. University of Michigan, Ann Arbor.

Siciliano, B. and Khatib, O. (2016). Robotics and the Handbook. In Siciliano, B. and Khatib, O. (eds.), *Springer Handbook of Robotics*. Springer, pp.1–6.

Skibniewski, M. (1996). Current Status and Key Issues for Construction Automation and Robotics in the USA. In *Proceedings of the 13th International Symposium on Automation and Robotics in Construction (ISARC'96)*, pp.51–57.

Skibniewski, M. (1986). *Engineering and Economic Analysis of Robotics Application Potential in Selected Construction Operations.* (Doctoral dissertation, Carnegie Mellon University).

Skibniewski, M. (1984). S*trategy for Engineering and Economic Analysis of Robotics Applications Potential in Selected Construction Work Operations*. (Ph.D. Thesis Proposal, Carnegie Mellon University).

Skibniewski, M. and Golparvar-Fard, M. (2015). Toward a Science of Autonomy for Physical Systems: Construction. A white paper prepared for the Computing Community Consortium (CCC) Committee of the Computing Research Association.

Skibniewski, M. and Russell, J. S. (1989). Robotic Applications to Construction. In *Cost Engineering*, *31*(6), pp.10–18.

Skibniewski, M., Derrington, P., and Hendrickson, C. (1986). Cost and Design Impact of Robotic Construction Finishing Work. In *CAD and Robotics in Architecture and Construction: Proceedings of the Joint International Conference at Marseilles*. pp.169–186.

Soden, R., Ribes, D., Avle, S., Fox, S. E., Sengers, P., Paudel, S., and Marathe, M. (2023). Historicism in/as CSCW Method: Research, Sensibilities, and Design. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*, pp.497–500.

Soden, R., Ribes, D., Avle, S., and Sutherland, W. (2021). Time for Historicism in CSCW: An Invitation. *Proceedings of the ACM on Human-Computer Interaction*, *5*(CSCW2), pp.1–18.

Soden, R., Ribes, D., Jack, M., Sutherland, W., Khovanskaya, V., Avle, S., Sengers, P., and Bødker, S. (2019). Fostering Historical Research in CSCW & HCI. In *Conference Companion Publication of the CSCW 2019 on Computer Supported Cooperative Work and Social Computing*, pp.517–522.

Son, H., Kim, C., Kim, H., Han, S. H., and Kim, M. K. (2010). Trend Analysis of Research and Development on Automation and Robotics Technology in the Construction Industry. *KSCE Journal of Civil Engineering*, *14*, pp.131–139.

Sorenson, J. (2019). Toward a Pragmatic and Social Engineering Ethics: Ethnography as Provocation. *Paladyn, Journal of Behavioral Robotics*, *10*(1), pp.207–218.

SRI International Youtube Channel (2017). Shakey the Robot: The First Robot to Embody Artificial Intelligence. [online] Available at: https://www.youtube.com/watch?v=7bsEN8mwUB8 [Accessed 28 September 2022].

Srihari, K. and Deisenroth, M. P. (1988). Robot Programming Languages — A State of the Art Survey. In *Robotics and Factories of the Future '87: Proceedings of the Second International Conference.* pp.625–635.

Stanford University InfoLab (1991). TAKE ME, I'M YOURS: The Autobiography of SAIL. [online] Available at: http://infolab.stanford.edu/pub/voy/museum/pictures/AIlab/SailFarewell.html [Accessed 7 September 2022].

Suchman, L. (2007). *Human-machine Reconfigurations: Plans and Situated Actions*. Cambridge University Press.

Suchman, L. (2002a). Located Accountabilities in Technology Production. *Scandinavian Journal of Information Systems*, *14*(2), pp.91–105.

Suchman, L. (2002b). Practice-based Design of Information Systems: Notes From the Hyperdeveloped World. *The Information Society*, *18*(2), pp.139–144.

Suchman, L. (1993). Technologies of Accountability: Of Lizards and Aeroplanes. In Button, G. (ed.), *Technology in Working Order: Studies of Work, Interaction, and Technology*. Routledge, pp.113–126.

Suchman, L. (1987). *Plans and Situated Actions: The Problem of Human-Machine Communication.* Cambridge University Press.

Suchman, L. (1983). Office Procedure as Practical Action: Models of Work and System Design. *ACM Transactions on Information Systems (TOIS)*, *1*(4), pp.320–328.

Suchman, L., Blomberg, J., Orr, J. E., and Trigg, R. (1999). Reconstructing Technologies as Social Practice. In *American Behavioral Scientist*, *43*(3), pp.392–408.

Sung, J. Y., Grinter, R. E., Christensen, H. I., and Guo, L. (2008). Housewives or Technophiles? Understanding Domestic Robot Owners. In *Proceedings of the 3rd ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.129–136.

The Defense Advanced Research Projects Agency (DARPA) (2023). ARPA is Born. [online] Available at: https://www.darpa.mil/about-us/timeline/dod-establishes-arpa [Accessed 17 June 2023].

The Encyclopaedia Britannica – Dennis, M. A. (2024). Marvin Minsky American Scientist. [online] Available at: https://www.britannica.com/biography/Marvin-Lee-Minsky [Accessed 17 June 2023].

The Encyclopaedia Britannica (2023). Westinghouse Electric Corporation American Company. [online] Available at: https://www.britannica.com/topic/Westinghouse-Electric-Corporation [Accessed 17 June 2023].

The Henry Ford (2022). Robot, First Unimate Robot Ever Installed on an Assembly Line, 1961. [online] Available at: https://www.thehenryford.org/collections-and-research/digital-collections/artifact/183434/#slide=gs-450962 [Accessed 5 September 2022].

The Institute of Electrical and Electronics Engineers (IEEE) (2020). What Is a Robot?. [online] Available at: https://robots.ieee.org/learn/what-is-a-robot/ [Accessed 17 September 2021].

The International Association for Automation and Robotics in Construction (IAARC) (2021). History of IAARC. [online] Available at: https://www.iaarc.org/about/history [Accessed 17 September 2021].

The Library of Congress (2023). The Westinghouse Electric and Manufacturing Company. [online] Available at: https://www.loc.gov/collections/films-of-westinghouse-works-1904/articles-and-essays/the-westinghouse-world/the-westinghouse-electric-and-manufacturing-company/ [Accessed 17 June 2023].

The Link – Togyer, J. (2012). RI Co-founder Tom Murrin: 1929-2012. [online] Available at: https://www.cs.cmu.edu/link/ri-co-founder-tom-murrin-1929-2012 [Accessed 10 September 2022].

The Manufacturing Futures Institute (MFI) (2022). Mill 19. [online] Available at: https://engineering.cmu.edu/mfi/mill19/index.html [Accessed 10 November 2022].

The Mobile Robot Laboratory Staff (1984). Towards Autonomous Vehicles. In *The Robotics Institute 1984 Annual Research Review*. The Robotics Institute, Carnegie Mellon University.

The Museum of Applied Arts & Sciences Australia (The Powerhouse) – Hutchison, A. (2019). Cincinnati Milacron T3 Industrial Robot. [online] Available at: https://collection.powerhouse.com.au/object/574553 [Accessed 17 June 2024].

The National Commission on Technology, Automation, and Economic Progress (1966). *Report of the National Commission on Technology, Automation, and Economic Progress: Technology and the American Economy, Volume 1*. U.S. Government Press Office.

The National Institute of Advanced Science and Technology (AIST) (2018). Development of a Humanoid Robot Prototype, HRP-5P, Capable of Heavy Labor. [online] Available at: https://www.aist.go.jp/aist_e/list/latest_research/2018/20181116/en20181116.html [Accessed 28 December 2023].

The National Research Council (1999). Developments in Artificial Intelligence. In The National Research Council, *Funding a Revolution: Government Support for Computing Research*. National Academies Press, pp.198–225.

The New York Times (2010). Stanford Cart. [online] Available at: https://www.nytimes.com/video/science/1247468063802/stanford-cart.html [Accessed 28 September 2022].

The New York Times – Ravo, N. (1998). Richard M. Cyert Dies at 77; President of Carnegie Mellon. [online] Available at: https://www.nytimes.com/1998/10/10/us/richard-m-cyert-dies-at-77-president-of-carnegie-mellon.html [Accessed 15 September 2022].

The New York Times (1982). Sale of Unimation to Westinghouse. [online] Available at: https://www.nytimes.com/1982/12/07/business/sale-of-unimation-to-westinghouse.html [Accessed 5 September 2022].

The New York Times – Stokes, H. S. (1982). Japan's Love Affair with the Robot. [online] Available at: https://www.nytimes.com/1982/01/10/magazine/japan-s-love-affair-with-the-robot.html [Accessed 17 June 2023].

The New York Times – Farnsworth, C. H. (1982). Japan-U.S. Trade Gap $13.4 Billion. [online] Available at: https://www.nytimes.com/1982/01/19/business/japan-us-trade-gap-13.4-billion.html [Accessed 17 June 2023].

The New York Times – Lohr, S. (1981). New in Japan: The Manless Factory. [online] Available at: https://www.nytimes.com/1981/12/13/business/new-in-japan-the-manless-factory.html [Accessed 17 June 2023].

The New York Times – Tharp, M. (1981). U.S. and Japan Agree on Ceilings for Car Shipments Through 1983. [online] Available at: https://www.nytimes.com/1981/05/01/business/us-and-japan-agree-on-ceilings-for-car-shipments-through-1983.html [Accessed 17 June 2023].

The Office of Naval Research (2023). All ONR-sponsored Nobel Laureates. [online] Available at: https://www.nre.navy.mil/about-onr/history/nobels [Accessed 17 June 2023].

The Office of Naval Research – Babb, C. E. (2022). Records Overview for Researchers. [online] Available at: https://www.nre.navy.mil/about-onr/history/history-research-guide [Accessed 10 September 2022].

The Pandas Development Team (2024). Pandas. [online] Available at: https://pandas.pydata.org/ [accessed 3 July 2022].

The Piper – Gerson, B. (2020). Demolition to Pave Way for New Scaife Hall. [online] Available at: https://www.cmu.edu/piper/news/archives/2020/october/scaife-hall.html#:~:text=Demolition%20of%20Scaife%20Hall%20is,of%20the%202023%20fall%20semester [Accessed 3 November 2022].

The RoBot Intelligence Group (CMU BIG). Social Navigation Simulator (SNS). [online] Available at: https://github.com/cmubig/Social-Navigation-Simulator [Accessed 15 July 2024].

Thibadeau, R. (1984). Automatic Visual Inspection as Skilled Perception. In *The Robotics Institute 1984 Annual Research Review*. The Robotics Institute, Carnegie Mellon University.

Thunberg, S. and Ziemke, T. (2020). Are People Ready for Social Robots in Public Spaces?. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pp.482–484.

Time Magazine (1961). Business: The Automation Jobless. [online] Available at: https://content.time.com/time/subscriber/article/0,33009,828815-1,00.html [Accessed 3 June 2023].

Tombesi, P. (2010). On the Cultural Separation of Design Labor. In Deamer and Bernstein (eds.), *Building (in) the Future: Recasting Labor in Architecture*. Princeton Architectural Press, pp.117–136.

Tory, M., Staub-French, S., Po, B. A., and Wu, F. (2008). Physical and Digital Artifact-mediated Coordination in Building Design. *Computer Supported Cooperative Work (CSCW)*, *17*, pp.311–351.

Tractica (2019). Construction Demolition Robots (cited in Business Wire (2019)). [online] Available at: https://www.businesswire.com/news/home/20190507005059/en/Construction-Robotics-Market-to-Reach-226-Million-Worldwide-by-2025-According-to-Tractica [Accessed 11 August 2024].

Trautman, P. and Krause, A. (2010). Unfreezing the Robot: Navigation in Dense, Interacting Crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp.797–803.

TurtleBot (2023). TurtleBot. [online] Available at: https://www.turtlebot.com/about/ [Accessed 28 December 2023].

Ueno, T., Maeda, J., Yoshida, T., and Suzuki, S. (1986). Construction Robots for Site Automation. In *CAD and Robotics in Architecture and Construction: Proceedings of the Joint International Conference at Marseilles*. pp.259–268.

Ulset, M. O., Broadbent, E., and Eriksen, T. H. (2022). Automated Care in New Zealand. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp.959–962.

Venkatesh, V. and Davis, F. D. (2000). A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies. *Management Science*, *46*(2), pp.186–204.

Vertesi, J. (2019). From Affordances to Accomplishments: Powerpoint and Excel at NASA. In Vertesi, J. and Ribes, D. (eds.), *DigitalSTS: A Field Guide for Science & Technology Studies*. Princeton University Press, pp.369–392.

Vertesi, J. and Ribes, D. (eds.) (2019). *DigitalSTS: A Field Guide for Science & Technology Studies*. Princeton University Press.

Vora, C., Aryal, A., Willoughby, S., and Wang, C. (2024). Investigating Stakeholder Perception and Developing a Decision Framework for Robot Adoption in Construction. *Journal of Construction Engineering and Management*, *150*(4), p.04024012.

Wainer, J. and Barsottini, C. (2007). Empirical Research in CSCW—A Review of the ACM/CSCW Conferences From 1998 to 2004. *Journal of the Brazilian Computer Society*, *13*, pp.27–35.

Wallace, J. (2021). Getting Collaborative Robots to Work: A Study of Ethics Emerging During the Implementation of Cobots. In *Paladyn Journal of Behavioral Robotics*, *12*(1), pp.299–309.

Wallace, J. R., Oji, S., and Anslow, C. (2017). Technologies, Methods, and Values: Changes in Empirical Research at CSCW 1990-2015. In *Proceedings of the ACM on Human-Computer Interaction*, *1*(CSCW), pp.106:1–106:18.

Wallén, J. (2008). *The History of the Industrial Robot*. Linköping University Electronic Press.

Warnecke, H. J. and Schraft, R. D. (1982). *Industrial Robots: Application Experience*. IFS Publications.

Warszawski, A. (1985). Economic Implications of Robotics in Building. In *Building and Environment*, *20*(2), pp.73–81.

Warszawski, A. (1984). Application of Robotics to Building Construction. In *Proceedings of the 1st International Symposium on Automation and Robotics in Construction (ISARC)*. International Association for Automation and Robotics in Construction (IAARC).

Warszawski, A. and Navon, R. (1998). Implementation of Robotics in Building: Current Status and Future Prospects. *Journal of Construction Engineering and Management*, *124*(1), pp.31–41.

Warszawski, A. and Sangrey, D. (1985). Robotics in Building Construction. In *Journal of Construction Engineering Management*, *111*(3), pp.260–280.

Wartzman, R. (2017). The First Time America Freaked Out Over Automation. [online] Available at: https://www.politico.com/magazine/story/2017/05/30/rick-wartzman-book-excerpt-automation-donald-trump-215207/ [Accessed 3 June 2023].

Whittaker, W. L. (2021). Interviewed by Emek Erdolu [Zoom]. 13 February.

Whittaker, W. L. (1986). Construction Robotics: A Perspective. In *CAD and Robotics in Architecture and Construction: Proceedings of the Joint International Conference at Marseilles*. pp.105–112.

Whittaker, W. L. and Bandari, E. (1986). A Framework for Integrating Multiple Construction Robots. In *CAD and Robotics in Architecture and Construction: Proceedings of the Joint International Conference at Marseilles*. pp.159–164.

Whittaker, W. L., Crowley, J. L., Oppenheim, I. J., Bares, J., Wood, T. A., Berman, S., and Lee, K. (1985). Mine Mapping by a Robot with Acoustic Sensors. In *Proceedings of the Second International Conference on Robotics in Construction*.

Wilson, P. (1991). Computer Supported Cooperative Work (CSCW): An Overview. In *Digital Creativity (Intelligent Tutoring Media)*, *1*(3), pp.103–116.

Winkle, K., Caleb-Solly, P., Turton, A., and Bremner, P. (2018). Social Robots for Engagement in Rehabilitative Therapies: Design Implications From a Study with Therapists. In *2018 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI),* pp.289–297.

Winograd, T. (1973). A Procedural Model of Language Understanding. In Schank, R. C. and Colby, K. M. (eds.), *Computer Models of Thought and Language*. WH Freeman & Co., pp.152–186.

Winograd, T. (1971). *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*. (Doctoral dissertation, Massachusetts Institute of Technology).

Wise, M., Ferguson, M., King, D., Diehr, E., and Dymesich, D. (2016). Fetch and Freight: Standard Platforms for Service Robot Applications. In *Workshop on Autonomous Mobile Service Robots*, pp.1–6.

Woodbury, R. F. and Oppenheim, I. J. (1988). An Approach to Geometric Reasoning in Robotics. *IEEE Transactions on Aerospace and Electronic Systems*, *24*(5), pp.630–646.

Wright, P. K. (1983). A Manufacturing Hand. In *The Robotics Institute 1983 Annual Research Review*. Carnegie Mellon University, pp.21–31.

Wu, Y. (2024). *A Human-Centered Reinforcement Learning Driven Robotic Framework for Assisting Construction Workers.* (Doctoral dissertation, Carnegie Mellon University).

Xiao, B., Chen, C., and Yin, X. (2022). Recent Advancements of Robotics in Construction. *Automation in Construction*, *144*, p.104591.

Yablonina, M., Nicolas K. K., and Menges A. (2020). Designing [with] Machines: Task- and Site-Specifc Robotic Teams for In-Situ Architectural Making. In *Proceedings of the 40th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA)*, pp.330–339.

Yasuoka, M. (2015). Collaboration Across Professional Boundaries—The Emergence of Interpretation Drift and the Collective Creation of Project Jargon. *Computer Supported Cooperative Work (CSCW)*, *24*(4), pp.253–276.

Yasuoka, M. (2009). *Bridging and Breakdowns: Using Computational Artifacts Across Social Worlds*. (Doctoral dissertation, University of Copenhagen).

Yuan, S., Coghlan, S., Lederman, R., and Waycott, J. (2022). Social Robots in Aged Care: Care Staff Experiences and Perspectives on Robot Benefits and Challenges. In *Proceedings of the ACM on Human-Computer Interaction*, *6*(CSCW2), pp.1–23.

Zamalloa, I., Kojcev, R., Hernández, A., Muguruza, I., Usategui, L., Bilbao, A., and Mayoral, V. (2017). Dissecting Robotics—Historical Overview and Future Perspectives. Preprint available on arXiv.org (arXiv:1704.08617).

Zhang, J. and Singh, S. (2014). LOAM: Lidar Odometry and Mapping in Real-time. In *Robotics: Science and Systems* (Vol.2, No.9), pp.1–9.

Zhu, A., Pauwels, P., and De Vries, B. (2021). Smart Component-oriented Method of Construction Robot Coordination for Prefabricated Housing. In *Automation in Construction*, *129*, pp.103778.

Image references

Chapter 2

Figure 1.
Image showing Unimate at the General Motors factory (1961-1962) (resized).
Image source: The Henry Ford.
Reused with permission from The Henry Ford.

Figure 2.
Still from the video “Butterfinger (AI Film #9)” (1967).
Video source: MIT CSAIL www.csail.mit.edu/node/6649.
Reused with permission from MIT CSAIL.

Figure 3.
Victor Scheinman next to the Stanford Hydraulic Arm (c.1968) (resized).
Image source: Stanford University InfoLab infolab.stanford.edu/pub/voy/museum/pictures/display/1-Robot.htm.

Figure 4.
Page from “The Robotics Institute 1983 Annual Research Review” showing Raj Reddy (1983).
Image source: The Robotics Institute 1983 Annual Research Review.
Copyrighted material, reused under non-commercial permitted use clause.

Figure 5.
“Figure 2: Comparison of Production Technologies” (1981).
Image credit: Robert Ayres and Steven Miller.

Image source: © Steven Miller.
Reused with permission from Steven Miller.

Figure 6.
“Figure 4: Three Levels of Control with The Cell Supervisor” (1983).
Image credit: David Bourne and Mark Fox.
Image source: The Robotics Institute 1983 Annual Research Review.
Copyrighted material, reused under non-commercial permitted use clause.

“Figure 5: The Initial Test Site for TRANSCELL” (1983).
Image credit: David Bourne and Mark Fox.
Image source: The Robotics Institute 1983 Annual Research Review.
Copyrighted material, reused under non-commercial permitted use clause.

Figure 7.
Photo of Shakey in a lab setup (1973).
Image source: © SRI International.
Reused with permission from SRI International.

Figure 8.
Stills from the video “RRV: Three Mile Island Rover Overview: 1984 CMU Robotics Institute Archive” (1984).
Video source: Carnegie Mellon University Robotics Institute, Field Robotics Center.
Reused with permission from Carnegie Mellon University Robotics Institute Video Archive.

Figure 9.
Still from the video “Terregator - CMU - Field Robotics Circa 1983” (c.1983).
Video source: Carnegie Mellon University Robotics Institute, Field Robotics Center.
Reused with permission from Carnegie Mellon University Robotics Institute Video Archive.

Figure 10.
“Figure 1: System Architecture” (1986).
Image credit: William Whittaker and Esfandiar Bandari.
Image source: CAD and Robotics in Architecture and Construction, Proceedings of the Joint International Conference at Marseilles. © 1986 Springer Nature BV.
Reproduced with permission of Springer Nature BV via Copyright Clearance Center.

Figure 11.
“Table 1.–Basic Activities in Building Construction” (1985).
Image credit: Abraham Warszawski and Dwight Sangrey.
Image source: © The American Society of Civil Engineers (ASCE).
Reused with permission from the American Society of Civil Engineers (ASCE).

Figure 12.
“Figure 3.–Construction Robots: (a) Assembly Robot; (b) General Purpose Robot; (c) Floor Finishing Robot; (d) Exterior Wall Robot” (1985).
Image credit: Abraham Warszawski and Dwight Sangrey.
Image source: © The American Society of Civil Engineers (ASCE).
Reused with permission from the American Society of Civil Engineers (ASCE).

Figure 13.
“Figure 3-2: General setup of the Sandblasting System” (1986).
Image credit: Miroslaw Skibniewski.
Image source: © Miroslaw Skibniewski.
Reused with permission from Miroslaw Skibniewski.

Figure 14.
“Figure 2-1: A Modular Approach to Robotic System Design” (1986)
Image credit: Miroslaw Skibniewski.
Image source: © Miroslaw Skibniewski.
Reused with permission from Miroslaw Skibniewski.

Figure 15.
“Table 1 – Example Construction Robotic Prototypes Developed in the United States” (1989).
Image credit: Miroslaw Skibniewski and Jeffrey Russell.
Image source: © 1989 by AACE International, all rights reserved.
Reused with permission from AACE International, 726 East Park Ave., #180, Fairmont, Wv 26554.
Email: info@aacei.org; Phone 304.296.8444; website: web.aacei.org.

Chapter 4

Figure 1.
Photo of Hazelwood in the early 1900s (resized).
Image source: Almono Limited Partnership www.almono.org/thehistory.

Photo of Hazelwood in post 2019 (resized).
Image source: Carnegie Mellon University
www.cmu.edu/news/stories/archives/2022/september/global-energy-leaders-view-innovation-at-mill-19.

Figure 2.
Photo of the Mill 19 facility (resized).
Image source: Carnegie Mellon University
www.cmu.edu/news/stories/archives/2019/september/mill-19-grand-opening.html.

Figure 3.
3D model drawing of the Mill 19 ground floor.
Image source: The Manufacturing Futures Institute (MFI) engineering.cmu.edu/mfi/mill19/index.html.

Figure 15.
Photo of Husky robot with sensors (2023).
Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

3D model drawing of the final container design (2023).
Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

Figure 16.
3D model drawing of LiDAR vision cone (2022).

Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

3D point cloud reconstruction of lab space (2022).
Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

Figure 17.
Diagram showing the two-step reinforcement learning approach (2022).
Image credit: ReAC Research Group / Yuning Wu.

Figure 18.
Stills from the simulation (2023).
Image credit: ReAC Research Group / Yuning Wu.

Drawings showing the lab test setup (2023).
Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

Photos of the lab test setup (2023).
Image credit: ReAC Research Group / Yuning Wu and Jiaying Wei.

Chapter 5

Figure 11.
Screenshots of Basil user interface (2022).
Image source: Canvas and Formant.
Reused with permission from Canvas.

Chapter 6

Figure 7.
Screenshots of Basil user interface (2022).
Image source: Canvas and Formant.
Reused with permission from Canvas.

Appendix

Appendix A – Key Definitions: What Is Work, What Is a Robot, What Is Robotics?

The Google-Oxford Languages Dictionary (2024) defines the verb "work" as "be[ing] engaged in physical or mental activity in order to achieve a result." The word has a slightly more complex connotation in this dissertation, which requires going beyond its dictionary definition. Work is first a "polymorphous concept," as sociologist Kjeld Schmidt (2011) determined based on philosopher Gilbert Ryle's (1951) framing. As a polymorphous concept, the specific attributes of work are open to interpretation until the activities it contains are specified. For example, according to Ryle's (1951) description of attributes, work may be bodily, performed on materials or with an instrument, done recreationally or professionally for pay. Without any context, we cannot know which of these attributes the word holds. Second, work is an ambiguous word. As Schmidt (2011) exemplifies, it may mean "designated places of work activities" when we say, "coming to work," "task" when we say, "taking work home," or "the outcome of work activities" when we say, "collected works."

Often used as a synonym for work, "labor" is also defined as "physical or mental effort" as well as "the services performed by workers for wages" or "an economic group comprising those who do manual labor or work for wages" in the Merriam-Webster Dictionary (2024). Based on this definition, the concept of labor is synonymous with what Ryle (1951) described as professional or occupational work. In other words, among many forms of work, this definition identifies labor as paid forms of work that are performed professionally, such as construction work, but it excludes unpaid forms that are performed recreationally, such as cooking or gardening. However, sociologists of work argue that labor is also present in unpaid recreational tasks (Grint, 2005). In this sense, like the concept of work, labor can be paid or unpaid and is still synonymous with work in both cases. Like the word "work," the word "labor" is also ambiguous. Depending on the use, labor may refer to "activities and effort that make up work" as well as "a collective body of people who perform those activities."

Based on these specifications, I take work and labor as synonymous and use them interchangeably in this dissertation. When I use these terms, I mean physical or mental activities performed on materials or with instruments and professionally as part of an occupation for pay. This definition associates work or labor with the concept of "practice," which I also employ often in the dissertation. Drawing on Schmidt's (2011) definition of practice, I consider the physical or mental activities that constitute work or labor as "rule-governed," which means that one has to learn and eventually master skills and techniques to perform them. For example, construction work, which is the object of study in the dissertation, refers to the body of physical and mental activities regularly and professionally performed on materials or with instruments to execute construction tasks. Carrying out these activities requires training to learn the essential skills and techniques as well as practical experience to potentially master them over time.

What is a "robot," and what is "robotics"? Contrary to popular imagination, a robot is distinguished from any other machine based on its specific components and capabilities rather than its human-like nature. For instance, the Institute of Electrical and Electronics Engineers (IEEE) (2020) defines a robot as "an autonomous machine capable of sensing its environment, carrying out computations to make decisions, and performing actions in the real world." The IEEE's (2020) definition suggests that a robot possesses a combination of sensing, software, and mechanical systems that enable it to autonomously sense, make decisions, and perform actions in the physical environment. This basic description is helpful to emphasize technical components and capabilities over appearance in distinguishing a robot from other machines. Yet, to fully describe what a robot is, we may need a further understanding of the concept of autonomy as well as the capabilities that make a machine a robot.

Based on researcher Jenay Beer and colleagues' (2014) definition, autonomy is "the extent to which a robot can sense its environment, plan based on that environment, and act upon that environment with the intent of reaching some task-specific goal—either given to or created by the robot—without external control." Building on this definition, the autonomy of a robotic system is dependent on the external control involved in reaching the intended goal. The field of robotics defines autonomy relative to the degree of external control; in this sense, autonomy exists on a spectrum, which is formalized in, for instance, the scale proposed by roboticists Michael Goodrich and Alan Schultz (2007). According to this scale, autonomy can range from "direct control" on one end to "dynamic autonomy" on the other, depending on external control ranging from "teleoperation" to "peer-to-peer collaboration." Therefore, the word "autonomous" in the IEEE's (2020) definition implies levels, and robots can have varying levels of autonomy relative to levels of control.

Regarding the capabilities that characterize a robot, there is an overall alignment. A robot is a system that senses, computes—also sometimes referred to as "thinking"—and acts accordingly. However, nuances remain in the identification of "act" when discussing robots rather than other machines operated by a learning- or model-based sensing system. For instance, computer scientist and roboticist Rodney Brooks (2020) defines a robot as "some sort of device which has sensors, which senses the world, does some sort of computation, decides on an action, and then does that action based on the sensor input, which makes some change out in the world outside its body." In contrast, according to computer scientist Gill Pratt (2020), even a conventional thermostat is a robot because just like a robot that senses, thinks, and acts, a thermostat senses the temperature, thinks about whether it is hot or cold, and acts by increasing or decreasing it.

Brooks's (2020) definition reminds us of machines with mechanical parts that can move and directly manipulate things in a setting through motion, such as Honda Asimo (a humanoid for assistance), ABB 6700 (a large industrial arm), or Roomba (a robotic vacuum cleaner). Conversely, Pratt's (2020) definition reminds us of systems without any visible mechanical parts or motion that nonetheless have the capability to manipulate their conditions, such as Amazon Alexa (a virtual assistant technology). The on-site construction robots I describe in this dissertation fall closer to Brooks's (2020) definition. They are the type of machines that can make changes outside their bodies. Unlike Asimo, they are not necessarily human-like in appearance. Unlike a thermostat or Alexa, they contain an explicit mechanical system and, at times, a body that allows them to move and manipulate their surroundings with motion. To refer to these devices in this dissertation, I interchangeably use the terms "robotic technologies," "robotic systems," and "robots." To denote those machines developed for construction, I add the word "construction" to these terms.

Lastly, "robotics" can be used to refer to the scientific and technological research field that deals with the design, development, and application of robotic technologies, as I frequently do in the dissertation. Aside from the research field, I also refer to practices that conduct the same activities in the industry. Although some sources identify robotics as an inter-discipline between mechanical engineering, electrical engineering, and computer science, it is also commonly identified as a branch of artificial intelligence (AI). In this dissertation, I also define robotics as an inter-discipline that includes computing and AI as its sub-parts. Construction robotics is a part of this inter-discipline that focuses on the design, development, and applications of robots for construction. I consider the emergence of research and development of on-site robots to mark the beginnings of construction robotics. I regard the early 1980s, when these activities were institutionalized as an academic research field and the first prototypes were introduced to construction sites, as the birth of the field.[1]

1 Based on sources including Carl Haas and coauthors (1995); Kamel Saidi and coauthors (2016); Borja Garcia de Soto and Miroslaw Skibniewski (2020).

Appendix B – Extended Descriptions of Study Planning and Data Collection

Site: Carnegie Mellon University Robotics Institute, Pittsburgh, Pennsylvania

Data Collection

At the center of historical research was a two-part investigation of three major archives: the CMU-RI's "Publications" archive, an online repository that includes published works from the Robotics Institute since 1979; the CMU Library's Robot Archive, a mixed-media archive featuring annual research reports titled the *Robotics Institute Research Review* published between 1983 and 1994; and "CMU Robotics Institute," the CMU-RI's official YouTube channel, which contains video clips of robotic system demos and researcher interviews from the 1980s and 1990s.

The first part entailed site data collection to draw CMU's institutional context. It consisted of a high-level review of robotics publications between 1979 and 1989 from the Publications archive, the *Robotics Institute Research Review* issues between 1983 and 1989 from the Robot Archive, and the information narrated in the interviews and system demos on the YouTube channel. These sources, especially the *Robotics Institute Research Review* issues, along with online articles, newspapers, oral accounts helped trace the details of the Robotics Institute's launch, sponsors, and motivations. Reviewing the publications along with these other sources allowed me to identify and document the key figures active at the intersection of the Robotics Institute and the DCE. To position this local institutional context, I also referred to primary sources such as research reports, newspapers, digital collections, and secondary historical sources that informed me about the broader context of robotics and labor in the United States.

The second part involved a more extensive survey of publications on the work produced in CMU's institutional context. The review covered the CMU-RI and DCE publications between 1979 and 1989 as well as the 1980s *Robotics Institute Research Review* issues and allowed me to extract a set of key research studies and systems (of historical significance and relevance to the investigation) for close reading and examination. An initial longlist included a selection of items, offering a perspective or projection on work while describing a technical system or involving analyses of work in service of robotics. From this initial list, a subset of projects were closely examined and included in the discussions in the dissertation.

To obtain further details about CMU's institutional context and projects, I collected oral accounts from a subset of the key figures previously listed. Because I received fewer responses to my invitations from Robotics Institute members, I had to concentrate more on DCE members. Nonetheless, this worked out nicely because documentation about the DCE in archival or contemporary sources is much scarcer than for the Institute. I ended up interviewing three key members of the DCE with close connections to the Institute in the 1980s, one of which went on to become a researcher at the Institute in the mid-1980s. Additionally, I interviewed three other researchers who worked at the Robotics Institute at the time, one of which co-authored a report that I examined as a key material for the dissertation.[2] Other than these, I relied on two existing oral accounts from the founding director of the CMU-RI, which provided valuable details about the Institute's launch.[3] The interviews were conducted on Zoom, an online video conference application,

2 Except for the one with the co-author of the report I examined in the dissertation, these interviews did not provide any relevant information.

3 These two oral accounts from the founding director of the Robotics Institute are obtained by John Markoff (2018), a historian at the Computer History Museum, and Hansen Hsu (2021), a curator at the Computer History Museum Software History Center.

and transcribed via the app's audio transcription feature. In the interviews, I asked participants about the position of robotics in 1980s engineering research, general and individual motivations behind robotics research, how the CMU-RI came about, and how robotics research emerged at the DCE. I also asked the interviewees about the local context of the Institute and the DCE, their engagements with the Institute, and the robotics work they authored in the 1980s.

Key figures researched and listed as potential interviewees:

- Robert Ayres (EDRC / Robotics Institute, 1978–1983)
- Haruhiko Asada (Robotics Institute, 1983)
- David A. Bourne (Robotics Institute, Could not be located–Present)
- Mark R. Cutkosky (Robotics Institute, Could not be located–1985)
- Richard M. Cyert (Carnegie Mellon University President, 1972–1990)
- Steven J. Fenves (Department of Civil Engineering, 1972–1998)
- Mark S. Fox (Robotics Institute, 1979–1991)
- Takeo Kanade (Robotics Institute, 1980–Present)
- Raj Reddy (Robotics Institute, 1979–Present)
- Daniel R. Rehak (Department of Civil Engineering, 1981–Could not be located)
- Dwight Sangrey (Department of Civil Engineering, 1979–1985)
- Miroslaw Skibniewski (Department of Civil Engineering, 1982–1986)
- Robert Woodbury (School of Architecture / EDRC, 1982–1993)
- Paul K. Wright (Robotics Institute, 1979–1987)

Interviewed:

- Linda Argote (School of Business, 1979–Present)
- Chris Hendrickson (Department of Civil Engineering, 1978–Present)
- Steven Miller (EDRC / Robotics Institute, 1978–1983)
- Irving Oppenheim (Department of Civil Engineering / School of Architecture, 1972–Present)
- William Whittaker (Department of Civil Engineering / Robotics Institute, 1975–Present)

Items researched and listed for close reading (in chronological order):

- "Unimate" (1954, 1956)
    - A robotic arm by George Devol and Joseph Engelberger (Unimation Inc.)
- "The Stanford Cart" (1960)
    - A mobile robot by James L. Adams and later by various researchers at SAIL
- "The Orm" (1965)
    - A robotic arm used developed by various researchers at SAIL
- "The Hydraulic Arm" (1966)
    - A robotic arm used developed by various researchers at SAIL
- "The Hopkins Beast" (1965)
    - A mobile robot by John Hopkins University Applied Physics Laboratory
- "Shakey" (1966)
    - A mobile robot by SRI
- "The Tentacle Arm" (1968)
    - A robotic arm by Marvin Minsky and collaborators at MIT AI Lab
- "Programmable Universal Machine for Assembly (PUMA)" (1973, 1978)
    - A robotic arm by Victor Scheinman, Unimation, Inc., and General Motors (GM)
- "The Tomorrow Tool (T3)" (1974)
    - A robotic arm by Cincinnati Milacron
- "Direct-Drive Arm 1" (1981)

    - The first direct-drive robotic arm by Haruhiko Asada and Takeo Kanade
- Shimizu Site Robot (SSR) (1983)
    - A fireproofing robot by Shimizu Corporation
- “The Human Side of Robotics” (1983)
    - A research publication by Linda Argote and collaborators
- “Expert Systems in Civil Engineering, Construction, and Construction Robotics” (1984)
    - A research publication by Daniel R. Rehak and Steven J. Fenves
- “The Man-Machine Interface” (1984)
    - A research publication by Robert Ayres
- “Direct-Drive Arm 2” (1985)
    - The second direct-drive arm by Takeo Kanade
- “REX” (1985)
    - A robotic excavator by William Whittaker and collaborators
- “The Organizational Implications of Robotics” (1985)
    - A research publication by Linda Argote and Paul Goodman
- “Automation and Craftsmanship” (1988)
    - A research publication by David A. Bourne and Paul K. Wright
- “Manufacturing: Acquiring Craft Skills Through Dialogs” (1986)
    - A research publication by David A. Bourne
- “Application of Model Based Control to Robotic Manipulators” (1988)
    - A research publication by Lyman J. Petrosky and Irving J. Oppenheim
- “Model Based Manipulator Control” (1989)
    - A research publication by Lyman J. Petrosky and Irving J. Oppenheim
- “A Field Prototype of a Robotic Pavement Crack Sealing System”
    - A research publication by Carl Haas and collaborators (1992)

Examined via close reading:
- “Butterfinger” (1963)
    - A robotic arm used and further developed by various researchers at SAIL.
- “The Stanford Arm” (1969)
    - A robotic arm by Victor Scheinman and collaborators at SAIL
- “Final Report: The First Ten Years of Artificial Intelligence Research at Stanford” (1973)
    - A research publication by Lester Earnest and collaborators at SAIL
- “The Use of Sensory Feedback in a Programmable Assembly System” (1973)
    - A research publication by Robert Bolles and Richard Paul at SAIL
- “The Impacts of Industrial Robots” (1981)
    - A research publication by Robert Ayres and Steven Miller
- “A Comparative Study of Robot Languages” (1982)
    - A research publication by Kang Shin and Susan Bonner
- “Designing Programming Languages for Manufacturing Cells” (1982)
    - A research publication by David A. Bourne and Paul Fussell
- “Remote Reconnaissance Vehicle” (1982)
    - A mobile robot for nuclear recovery work by William Whittaker and collaborators
- “A Manufacturing Hand” (1983)
    - A robotic gripper by Paul K. Wright and Mark R. Cutkoski
- “Autonomous Manufacturing: Automating the Job Shop” (1983)
    - A research publication by David A. Bourne and Mark S. Fox
- “Terregator” (1983)
    - A mobile robot for indoor and outdoor work by William Whittaker and collaborators

- “TRANSCELL” (1983)
    - An operating system for a manufacturing cell by David A. Bourne and Mark S. Fox
- “Automatic Visual Inspection System” (1984)
    - A visual inspection system for manufacturing by Robert Thibadeau and collaborators
- “The Mobile Robot Laboratory robots – Pluto, Neptune, Uranus” (1984)
    - Mobile robots by Hans Moravec and collaborators
- “Robotics in Building Construction” (1985)
    - A research publication by Abraham Warszawski and Dwight A. Sangrey
- “A Framework for Integrating Multiple Construction Robots” (1986)
    - A research publication by William Whittaker and Esfandiar Bandari
- “Cost and Design Impact of Robotic Construction Finishing Work” (1986)
    - A research publication by Miroslaw Skibniewski and collaborators
- “Construction Robots for Site Automation” (1986)
    - A research publication by Tomohiro Ueno and collaborators at Shimizu Corporation
- “Construction Robotics: A Perspective” (1986)
    - A research publication by William Whittaker
- “Higher Order Languages for Robots” (1986)
    - A research publication by James Blaha and collaborators
- “Robot Programming Languages — A State of the Art Survey” (1988)
    - A research publication by Krishnaswami Srihari and Michael Deisenroth
- “Robotic Applications to Construction” (1989)
    - A research publication by Miroslav Skibniewski and Jeffrey S. Russell

Site: The Rethinking Automation in Construction Project, Pittsburgh, Pennsylvania

Planning and On-Site Improvisation

Designed in two layers, the ethnographic research in the ReAC project consisted of an ethnographic study of construction work and an ethnographic study of robotic system development, which the first study aided. The planning of the first study involved identifying data collection methods, communicating these methods to the general contractor and its subcontractors as participants, scheduling on-site data collection, and obtaining the required permissions. I initially identified four main data collection methods: on-site observations of building tasks supplemented by time-lapse recordings, photos, and site mappings; interviews with selected workers, questionnaires to capture worker interactions, and a technology design workshop between workers and the ReAC group. Together with the principal investigator and the member leading the system development, we presented these methods to the project and site managers to receive their feedback. Along with the principal investigator, I assured the managers that we would be sensitive to workers’ perceptions of the time-lapse recording, clearly communicate our intentions, and respect their privacy. We also had to ensure that we would be mindful of their time during shifts.

Scheduling was crucial in planning and was aimed at synchronizing the on-site data collection with the construction project’s schedule. In an ethnography of construction, the presence of participants on-site changes dynamically. Building trades are hired to work for only a specific period as subcontractors. Therefore, capturing a particular task by a trade is contingent on matching the schedule of data collection with the period of presence of that trade on site. For our study, this arrangement was made through an information session with our site contact and the project’s safety manager, in which they shared the construction schedule for ongoing and upcoming tasks. Based on this information, the data collection was tentatively scheduled for a three-and-a-half-month period

between mid-January and May. This period would allow the general contractor to complete a concrete deck on one of the floors for safe walking while allowing us to overlap with multiple trades and tasks. After the field methods and the schedule were in place, we applied for and obtained the CMU Institutional Review Board's (IRB) approval of the study protocol. As mandated by the general contractor, I, as the ReAC member setting foot on the site, also completed two construction safety training courses.

A couple of weeks was sufficient to learn that site conditions unavoidably cause improvisations in planning. For instance, weather conditions delayed the concrete pouring of the floor we selected to observe (named Level B), which postponed the start of my site visits for a month. Site conditions also shaped the implementation of the field methods. Three notable examples were the addition of what I began to call the "daily observation protocol," the elimination of expensive methods, and the mastering of communication with the worker participants. Level B was highly cluttered and crowded, with twenty to thirty-five workers at times, and required me to structure each observation day with a routine protocol. This protocol, which lists routine steps such as taking attendance of participants on the floor or planning the time-lapse camera's positions, was instrumental in daily decision-making regarding who to observe or where to position the camera. Due to extremely tight job schedules, it was not possible to reserve workers' time for more than one interview, so the questionnaires and design workshop were dropped from the list of methods. Lastly, the time I spent on the floor and the degree of rapport I built with the workers over time improved my reading of the right time to approach them.

Like for this first study, the planning for the second study focusing on the robotic system development work included the identification of data collection methods, scheduling, obtaining the required permissions, and communicating these methods to my ReAC colleagues as participants. I initially identified five main data collection methods: participant observations of system development activities, including weekly meetings, discussions, and hands-on work, the examination of the system modules and hardware that would constitute the prototype robot, interviews with selected colleagues, experience sampling via questionnaires to extract the group's collaborative work interactions, and collecting samples from group conversations via audio recording.[4] The questionnaires and collection of conversation samples were scheduled to match the duration of the first study, whereas the observations, examination of the system, and interviews were scheduled over a longer timeline. For IRB approval, this second study protocol was combined with the first one within a single application and approved as such. With no bureaucratic layer between me, as the researcher, and the participants, communicating the field methods was no challenge. After scheduling and obtaining IRB approval, an email stating the purpose and methods of the study as well as the terms of consent, privacy, and confidentiality was sufficient.

Data Collection

The data collection on Level B eventually relied on observations and interviews. Observations took place three days a week for seven weeks and were structured by the daily observation protocol I mentioned earlier. Each observation day lasted for five to six hours, mostly between 8:30 am and 1:30–2:00 pm, which is typically the end of the day on US construction sites. On each observation day, I concentrated on one or two building tasks, and I recorded the activities carried out via field notes, photos, and time-lapse recordings. I took notes with a pen in a letter-size notebook, and I captured photos with an Apple iPhone 12 camera, which was light and convenient to use with one

4 See Reed Larson and Mihaly Csikszentmihalyi (2014) for an overview of experience sampling method.

hand while carrying a notebook and small enough to keep in the pocket when not in use. For the time-lapse recordings, I used three Brinno TLC200Pro HDR cameras mounted on tripods and powered by Anker power banks.

Time-lapse recording on a construction site turned out to be more complicated than taking notes or photos and, therefore, demanded more work in terms of setup and planning. Camera mounts made it challenging to position the cameras straight, and it was not possible to install them on the columns, which would soon be thickened with fireproofing retardant. Instead, I mounted each camera in a plastic casing on a tripod, with the lens at a height of approximately six feet two inches. I set the cameras to record at one-second time-lapse intervals, one frame per second, and with the best image quality along with default exposure, saturation, contrast, and sharpness, which worked best under the lighting conditions of Level B. I determined the cameras' positions at the beginning of each day. This was based on a brief observation of ongoing tasks to select the task(s) of focus as well as the available room and unobstructed view needed to capture it. On some days, consulting the workers on the duration of their stay in an area or the direction they would advance helped me determine camera positions as well.

Over twenty-one days on Level B, I collected notes, photos, and time-lapse recordings of seven tasks, including cleaning/housekeeping, fireproofing, rebar tying, carpentry formwork, mechanical electrical plumbing (MEP) ceiling line fittings (electrical, plumbing, steam fitting, sheet metal, and fire sprinkler fitting), framing, and drywall installation. At varying intervals, I reviewed the notes on these tasks and synthesized the salient points. At the end of each week, I organized the photos capturing these tasks, workers, tools, and materials in folders labeled by each day. I sorted the time-lapse recordings by day, camera, and task(s) recorded. In addition to this data, I also recorded the ways Level B was spatially organized into work and material storage zones each day. I did this via hand sketches on a plan drawing, delineating the extent of each zone through color-coding by trade and complemented by photos. At the end of the data collection, I cross-checked these sketches with the photos and, accordingly, made them more precise as digital drawings. The last data set I gathered was the frequency range of worker interactions, which I obtained by sampling daily observations—instead of questionnaires—after realizing that the number of these interactions within and across trades was small enough to observe and divide into range clusters.

Over the eight weeks following the end of the observations, I interviewed a total of nine workers. Because the interviews had to fall within work hours, scheduling was the biggest challenge and extended this process to eight weeks. Each interview also required an extra layer of communication with the superintendent to make it work with the weekly deadlines. Except for one on Level B, the interviews took place at a CMU building next door to the project site, and they lasted one hour fifteen minutes to one hour thirty minutes. I audio-recorded the conversations via Sony ICD-UX570 voice recorder along with Apple iPhone 12 voice recorder as backup. I also simultaneously took notes of important points in a notebook, and I used the same notebook as a drawing pad for the interviewee and myself to communicate by drawing. Workers who were involved in cleaning/housekeeping, fireproofing, rebar tying, plumbing, steam fitting, sheet metal, fire sprinkler fitting, framing, and drywall installation on Level B volunteered to interview. The data I obtained through the interviews includes the workers' descriptions of their tasks, the materials and tools they used, and the ways they coordinated. The data also includes the workers' perspectives on robots for construction and on the potential robot support use cases the ReAC group was considering at the time.

The second study employed all the methods I outlined in the plan except for interviews. In parallel with the data collection on Level B, participant observations of ReAC work continued one day a

week for fifteen weeks. These observations covered the weekly core group meetings joined by the principal investigator and the two research leads, including myself, weekly group meetings joined by all members, biweekly site information sessions in which I shared the on-site data and building task descriptions with the group, and weekly informal discussions between the research lead in charge of the system development and myself. I recorded the observations of all these events via field notes with a pen in a letter-size notebook. Additionally, I audio-recorded the conversations during the general group meetings and site information sessions with an Apple iPhone 12 voice recorder. Eventually, colleagues' views expressed naturally in those work moments and captured in these notes and audio recordings replaced the formal interviews.

Over this period, I also collected experience samples to examine collaborative work interactions and conversation samples by audio recording group meetings and site information sessions, as noted earlier. The weekly questionnaires I collected from the members over twelve weeks provided the experience samples. These questionnaires recorded pairwise interactions in the group, the medium used for each interaction (instant messaging, online video conference, in-person, etc.), and the initiator, theme, and frequency of each interaction on paper. They were later digitalized and sorted by member and week. I used Zoom's audio transcription feature to convert the audio recordings to text. The purpose of collecting this textual data was to trace how the knowledge of construction work gained from the first ethnographic study infiltrated the group's technology design discussions.

After the data collection on the construction site was completed, the group entered a summer period, in which we focused on further analyzing the construction data and developing a robot support scenario and the prototype system in parallel. During that period, I had to shift my attention to the Canvas study in San Francisco, so the participatory observations became remote observations of the biweekly online meetings. At that point, the mobile robot base that the system would be deployed on and the sensor suite that would be integrated into that base were settled, and the group had already produced some documentation including presentations, reports, and preliminary demo videos. Based on these materials, I started the examination of the software modules and hardware by focusing on the ways the construction data informed their implementation. I also worked on further analyzing a building task (carpentry formwork) in which we saw more potential for robot use and building a scenario for potential robot support in that specific task. Once the on-site part of the Canvas study ended in late summer, I resumed in-person observations and examinations and continued them until the end of the project. The documentation produced by the group and field notes were the primary recording tools during this final period.

Site: Canvas Construction Robotics, San Francisco, California

Planning and On-Site Improvisation

Arranging an ethnographic study on a site in the tech industry outside Pittsburgh required diligent planning. I started by communicating the objectives of the study to my contacts at Canvas and learning from them about the basics of what Canvas does. There was a period of productive exchange with one of the contacts about what Canvas might offer as a site for my investigation. This was followed by outlining the data collection methods, negotiating the use of these methods with the company, obtaining the IRB's permission, and scheduling the on-site data collection. The methods I initially proposed were observations of system development activities, including daily meetings, tests, and construction job site deployment(s), interviews with selected team members who conduct these activities, experience sampling via questionnaires to document their collaborative work interactions, and sampling of team conversations from online communication applications.

Additionally, an examination of Basil's software and hardware components as well as sampling recordings of the robot's user interface and robot system logs were included to expand the boundaries of data collection.

Negotiating the use of these methods with Canvas was an important step towards finalizing the data collection protocol. I discussed the type of data that each of these methods entails in a series of conversations with the site contacts. Employees' time was naturally the site contacts' central concern. Therefore, I had to make sure that the observations, interviews, and questionnaires would not be too obtrusive and time consuming. The site contacts' other concern was the sampling of conversations. Particularly, they were not in favor of textual data collection on the specific application that Canvas teams use daily, on which they sometimes share sensitive information. At their request, I removed this method from the protocol. In terms of fieldwork schedule, three-, five-, and seven-month options were discussed. As the ReAC study's timeline became clearer, a three-month fieldwork in the summer seemed feasible. Lastly, communicating the intentions and methods of the study with the rest of the Canvas team and seeking their voluntary participation was critical for the fieldwork to happen. With the site contacts acting as mediators, the methods and terms of consent, privacy, and confidentiality were communicated to the teams via email. Additional time was then spent on obtaining the consent forms due to the great number of volunteering participants.

Similar to the fieldwork on Level B, the site required some improvisation. For instance, familiarizing myself with the daily operations, including the division of remote and in-person labor around the system as well as the team and individual activities, turned out to be much more intensive than I estimated. This was largely because of the operation's size and the teams being spread between the Canvas office and individual homes.[5] The operation's size and the hybrid work mode also required me to prepare a more site-driven plan, which was dynamically adjusted based on daily new observations. For instance, after initial observations, recurring online team meetings and tests stood out as the key activities to concentrate on. The ones I found useful to observe expanded in number over the weeks and gradually formed a weekly observation calendar. Similarly, participants' shifting and often overlapping schedules kept the interview schedules moving and demanded a lot of coordination on the fly. Finally, because the number of participants made it time consuming to collect questionnaires, the meeting calendar app—in which the company included me to follow the events—became an improvised data source for analyzing work interactions.[6]

Data Collection

The Canvas study involved data collection in two broad categories: (1) collection of data related to Basil's development and (2) collection of data related to the broader work organization around the system. It relied on observations, computational data collection, and interviews, all concentrating on the company's engineering sub-teams.

Depending on the teams' schedule of activities, I distributed the observations into different times of each workday for eleven weeks. Observations within the first few weeks aimed at documenting the participants and individual roles. They focused on a list of online all-team meetings and a few in-person work activities and tests. I supplemented these observations with fifteen-to-thirty-minute information sessions or shorter chats with selected participants to learn about each team's ongoing

5 I worked with a total of sixty participants, and sorted out the key activities that I observed from a very large selection in this study.

6 The company uses Google Calendar, a commonly used scheduling app by Google. The app's user interface includes a calendar, populated with meeting tabs that the user adds and invites other attendants, or alternatively gets invited to. Each tab holds the time, attendants, location, and other meeting information. More information on Google Calendar app is widely available online.

work. These observations and interactions helped determine a larger list of meetings, tests, and potential job site deployments to keep an eye on in the following weeks. Accordingly, weekly observations increased to three all-team meetings, nine team meetings, and four inter-team meetings about Canvas's new targeted spray system, on which I focused for this study. In addition, I attended and observed a robot operator training session and three standard tests focusing on system functionality, system evaluation, and toolification, along with a job site deployment. Participation in these events, and observing Basil in action in some of them, enabled me to learn about the system's hardware and software components as well. I recorded the observations via field notes on a laptop and, when appropriate, photos and occasional short video recordings captured by an Apple iPhone 12 camera.

In addition to the observational data collected in person, I obtained two types of computational data that allowed for additional asynchronous observations. As previously mentioned, online team meetings scheduled via the meeting calendar app make up a key portion of the weekly work at Canvas, which positioned the app as a useful source for collecting what I call "work interactions data." I manually created this data by reformatting ten weeks of meeting information into sets of nested dictionaries in Python. Each week is represented by a set that contains a dictionary for each participant (anonymized with a pseudonym). Each participant's dictionary contains a dictionary for each meeting they attended (with the name of the meeting obfuscated), and each meeting dictionary includes a list of all attendees the participant interacted with. This data collection and curation also helped in understanding and documenting the participants, individual roles, and ongoing work. The other computational data were screenshots from recordings of the Basil's user interface and the system's log data, which were sourced from a client software that Canvas uses to monitor its systems' performance.[7] Samples of this data extended the observation of worker–robot interactions into software space during a system evaluation test and a job site deployment.

Further, a total of twenty-one participants from the engineering team were interviewed over the last six weeks of the fieldwork. Like for the interviews with construction workers in the ReAC study, finding available time slots in these individuals' busy and, at times, shifting schedules was a major challenge. I conducted most interviews, which lasted between one and two hours, in person in the meeting rooms of the Canvas headquarters, and the remaining were conducted online on Zoom. All interviews were audio-recorded via Zoom's recording feature and transcribed using the app. In all interviews, I used a whiteboard as a large-scale notepad to take notes and, in the in-person ones, to jointly diagram with the participant. I recorded these notes and diagrams with photos. Depending on the participants' varying roles, the interviews captured different kinds of information. Part of the interview data includes the fundamentals of drywall finishing work and the high-level computational aspects of Basil as well as participants' descriptions of how drywall finishing knowledge informs the system's development. Other parts include participants' descriptions of their roles and collaboration in the technology development work around the system, along with their general perspectives on robots for construction.

7 See formant.io for more information. System's log data was extracted and provided to me in tabular format by Canvas.

Appendix C – Descriptions of Building Tasks (continuing from Chapter 3)

Mechanical Electrical Plumbing (MEP) Ceiling Line Installations

> *I may do the same thing every day… but no two days are the same, never… To say that you can have everything prefab just to go up and in, you are more than likely not gonna work. (Fire Sprinkler Fitter 1)*
>
> *In this building, [given] how big our stuff is, if I just start in this corner and keep going, these other guys will cut my tires. Superintendent will kick me off the job. (Sheet Metal Worker 1)*

Throughout my time on Level B, MEP efforts were dedicated to installing ventilation, cooling, heating, plumbing, fire sprinkler, and electrical system lines on ceilings, which are embedded in BIM models or fabrication shop drawings. These lines are composed of specific material units, namely, ducts, steam pipes, plumbing pipes, fire sprinkler pipes, and electrical conduits. The installation of these units, also referred to as "fittings," is divided between specialized crews—sheet metal workers, steam fitters, plumbers, fire sprinkler fitters, and electricians, respectively. These crews first came to my attention when steam fitters, sheet metal workers, plumbers, and electricians took over the northeast section after fireproofing. A few workers and material stacks summarized their presence on those days. Stacks were imported to the floor similarly to how the garbage was exported, with cranes through the drone test arena. Stuff that was coming in was diverse: chilled water, hot water, steam pipes for heating and cooling; conduits for casing wiring; drainage, drinking water, lab air pipes for plumbing; and pipes for fire sprinklers, with lengths of up to twenty feet and diameters ranging from half an inch to ten inches, representing a mix of materials such as steel, PVC, copper, and cast iron. Materials also included bend-iron ducts, from main parts measuring up to fifty-six inches and a quarter to customized elbow turns and transitions. This operation, portrayed to me with words like "worst," "pain," or "tough" by workers, was initially infrequent but gradually increased as the task shifted gears. Soon, the floor was dominated by MEP workers and materials. As the task began to dominate, I started to pay close attention to MEPs and continued for the following two weeks.

As one of the largest in the building, Level B's ceiling comprised longer lines. For MEPs, this meant a greater quantity of the specialized hardware called the "hanger," made up of a threaded steel rod that is inserted into the ceiling and a bracelet-like hanger head attached to the rod to hold lines. More hangers meant laying out and drilling more pre-drafted points on the ceiling for rod insertion. Most of these points, which were spaced "anywhere from four to five feet to every ten feet," as Steam Fitter 1 noted, were drilled through Level A above before my visits to the site. Subsequently, MEPs dropped the rods through the corrugated deck separating Level A from Level B. The rest of this activity and the assembly of hangers with rods continued concurrently with line fittings on Level B. Few MEPs did layout work while I was present, but I managed to capture one instance, combined with hanger installation by Plumbers 1 and 7 for a drainage pipe close to the carpenters' formwork area (Figure 1). Plumber 1 started the task. Wearing a headlight, he was on his knees drawing the pipe's central axis on the ground with a pencil and a level he used as a ruler (Figure 1-A). Normally, the crew finds this axis by "pulling it off of" the closest north-south and east-west column lines, as Plumber 4 later explained in an interview. In this case, however, Plumber 1's main reference was a segment of the pipe installed previously. Once the axis was set, Plumber 1 grabbed a small device resembling a tape measure and positioned it on the penciled line. When turned on, the device projected a laser beam onto the ground, which extended the penciled line further. Along the laser

line, Plumber 1 took a measurement from the last hanger carrying the installed segment to find the next two hanger points. With one knee down, he penciled down the points. A puddle covered the rest of the line he had to measure. Using a squeegee, he swept the water away until the ground was suitable for penciling. He extended his tape measure from the first point, fixed its ends by bending down, and marked the next point on the wet concrete (Figure 1-B). He took the last measurement from this new point and marked the last point before taking a short break.

A few housekeeping steps followed after the break. Plumber 1 moved a garbage bin, a box, and one of the tool carts away. He picked his line laser up from the ground and put it back on the cart. He still had to wrestle with the puddle, so he did a few more rounds of sweeping in several directions. I soon realized that this preparation was to open up space for the lift, which he planned to maneuver along the axis. It was time for him to be both the ground and up guy simultaneously. He brought another laser device and positioned it on the first couple of hanger points he had penciled down. The device shot a beam upwards and projected the point with a red dot onto the corrugated deck. Plumber 1 looked up to check for any alignment errors. He performed the same precheck for the second and third points. After a short coordination chat with Plumber 7, who was hanging a pipe with Plumber 5 in the same area, he set the laser back on the first hanger point, climbed onto the lift, and positioned the vehicle right next to the laser. He elevated himself closer to the red dot on the deck, took a quick measurement, and came back down. He brought a steel strut and balanced this thin, long, rectangular stick on two tripod-like pipe stands. With a power tool called the “portable band saw,” he cut a small piece at the length of the measurement he took (Figure 1-C). Back on the lift with the piece, he positioned it over the red dot, located where the deck made a low section on its uniform, low-high-ridged profile. The sides of this low section were high sections, so he drilled a point on each of those and fixed the strut that now spanned the low section (Figure 1-D). I thought that there must be a constraint from the upper level because I saw workers drill and insert a rod directly into the red dot before. The answer came from Plumber 1 after a few more rounds of fixing struts for the other hanger points: it was because of the electrical conduits encased in the concrete deck above.

After the struts, Plumber 7 took over for the hanger installation. Following a short coordination chat covering the strut locations, rod thickness, and hanger specification for the cast iron drainage pipe, he hopped onto the lift with a level and elevated it until the installed portion of the pipe was at shoulder height. He inserted the level’s end into the pipe’s circular mouth and held the rest of it straight by stabilizing the opposite end with his left hand. Balancing a tape measure in his right hand, he measured the distance from the level to the strut (Figure 1-E). Meanwhile, Plumber 1 was giving me a live presentation. “He’s getting the length of the rod he needs,” he said, “and he used the pipe to find the next hanger [elevation].” When there is no pipe to help determine the elevation, workers find it by measuring up from a “benchmark elevation.” Later, Plumber 4 taught me the rest of the calculation: adding “[approximately] three-quarters of an inch” to reserve space for fastening along with a short length to each rod to create a gentle drainage slope. With the rod length in mind, Plumber 7 stepped down and grabbed a rod he had previously cut and attached it with a “clevis hanger,” as specified for drainage pipes. He double-checked that the length matched his measurement, trimmed it down with the portable band saw, and went back up with the sized piece. He bolted the piece into the strut with his hands and a hand tool called “channel lock pliers” (Figure 1-F). He lowered the lift, positioned the level between the previously installed pipe and the newly installed hanger, and checked the slope (Figure 1-G). After touching base with Plumber 1, he put the level on the hanger to find the next elevation and measured upward to determine the next hanger size. Accordingly, he prepared two new hangers, went back up, and installed them. He again checked the slope. For the next hanger, he performed the same routine. The spacing between this strut and the last one was longer than his level, so he instead installed the hanger first and then



Figure 1. Key moments of Plumbers 1 and 7's workflow on Level B

placed a long steel strut to see and iterate the slope. The crew left the floor after installing this hanger.

This hanger installation by Plumber 7 is only one sample. Depending on the line, variations on the same activity can entail greater complexity. For instance, the installation may involve more than one rod thickness, as Steam Fitter 1 taught me. On Level B, four sizes were used, including “three-eighths [of an inch of diameter] on the smallest [pipes],” “three-fourths on the bigger pipes,” and “one-halves and five-eighths in between.” A particular hardware may demand more than one rod per hanger, as I saw when the sheet metal crew adjoined every dropped rod for ducts with a second piece via additional hardware called the “spring isolator” for vibration isolation. This makes the calculation of the rod length more tedious than that performed by Plumber 7 above. Alternatively, the installation may require different types of hangers. For instance, other than the “clevis hanger,” I observed “racks,” that is, hangers that plumbers assemble with two rods and a strut for four or five parallel drinking water pipes to sit on. As Plumber 4 told me, plumbers also use racks for pipes running under ducts when “the only way to hang up a pipe below” is to span across. Conduits are structurally similar to drinking water lines, so electricians almost always use this type of hanger. Sheet metal workers call this hanger a “trapeze,” and they also use it for rectangular ducts that are thirty-six inches or wider. Lastly, depending on the type of line, hanger modification may be needed. For instance, for lines that require insulation, MEP uses oversized hangers, as I learned in interviews. The hangers Plumber 7 installed fit because drainage pipes do not necessitate insulation. However, for drinking water pipes, plumbers use oversized hangers. Similarly, steam fitters oversize their clevis and roller hangers to fit an insulation layer around their chilled water and steam pipes.

MEP hanger installations were followed by line fittings on Level B. I captured every MEP crew performing fitting, but the one by Steam Fitters 1, 2, and 4 on the middle section of the floor demonstrated every step of the activity (Figure 2). To prepare for the work, Steam Fitter 2 brought a scissor lift and a duct jack, a mechanical lift with pallets that MEP workers use to raise heavy items, while Steam Fitter 1 swept the floor, positioned the jack, and brought a ladder. The workers’ main stack was close to the ceiling portion to install. They walked and picked up a black steel pipe from the stack. With their shoulders under each end of the nearly twenty-foot piece, they carried it and carefully placed its center on the jack’s base (Figure 2-A). The section they would install was wide open, without any overlap with other MEP lines, so it seemed that they would use the entire piece with no need for sizing. After helping to position the pipe on the jack, Steam Fitter 2 jumped onto the lift. Steam Fitter 1 stayed with the jack and started to spin its handles, which lifted the pipe off the ground. Steam Fitter 2 followed the pipe raising by elevating his lift. His hands finally met the end of the pipe at a height of approximately twelve feet, a few inches from the previously installed clevis hanger on which the pipe would sit (Figure 2-B). Steam Fitter 2 took the hanger off the rod and pulled out the hardware’s bolt, which tightened its hold over the pipes. Meanwhile, Steam Fitter 1 locked the jack, pulled the ladder to the other end of the pipe, and climbed up. He performed the same steps with the hanger on his side (Figure 2-C). He then went down with the hanger in his hand and modified the jack’s position to place the pipe. Once the pipe was in the right spot, Steam Fitter 2 slid the hanger on it, put the bolt back on, and tightened it. Once he had screwed the hanger back on the rod, one side of the pipe lay on the hanger. A few seconds later, Steam Fitter 1 was done with the same steps on his side. As he lowered the duct jack, I saw that the pipe was hanging in full.

This pipe was a piece in the third line of a four-way system, cutting across the entire floor in sharp ninety-degree turns. Next would be adding to the fourth line, so Steam Fitter 1 moved the duct jack under the hangers of that line. With Steam Fitter 2, he walked to the stack and brought a full twenty-foot piece. As the workers approached with the pipe on their shoulders, Steam Fitter 4 joined them

with a lift (Figure 2-D). After placing the pipe on the jack, Steam Fitter 2 climbed back on his lift, and Steam Fitter 1 remained to steer the jack. Steam Fitters 2 and 4 moved their lifts, and each parked under a hanger, which would hold the pipe on opposite ends (Figure 2-E). Once all were set, Steam Fitter 1 started spinning and brought the pipe to the elevation of the workers on lifts. The lift team first balanced the pipe with their hands then elevated their lifts at a perfectly synchronous pace with the pipe going up. Once the pipe was at the proper elevation, they repeated the hanging routine. This time, instead of removing the entire hanger, they separated the hardware's U-shaped bottom part from the bolt on top and reassembled the two parts around the pipe (Figure 2-F). After this coordinated effort, Steam Fitter 4 left to work on another portion. Steam Fitter 1 moved the duct jack and ladder out of the way, and the space he opened up allowed Steam Fitter 2 to fill Steam Fitter 4's previous spot. As Steam Fitter 2 arrived, he received two elbow-style couplings from Steam Fitter 1, who picked them up from a box next to the pipe stack. He went up, cleaned the end of the pipe the crew had just fitted, and fastened a coupling that embraced and pressed into its end like a handcuff. Then, he connected the elbow to the coupling. (Figure 2-G). He fixed the other elbow to the second pipe and stepped down. These elbows, which were the last part of the crew's work in this area, would connect the pipes they installed to the next ones on the line.

Similar to hanger installation, variations of line fitting may involve additional complexity. For instance, the location of the activity may dictate material transfer. Within the steam fitting workflow I captured, part of the ceiling to be fitted was in close proximity to where the crew's pipes were stored. In many other instances, MEP workers carried batches from their main stack with carts and set temporary material stations for a particular area. Further, fitting in tight areas with different lines crossing may involve more material sizing. The steam fitter crew worked on a wide-open section, so they could install full twenty-foot segments. In other cases, ground guys often cut pipes into smaller lengths with specialized tools, such as the "four-wheel cutters" or the "portable band saw" for black steel pipes, the "reciprocating saw" for PVC pipes, and the "chain snapper" for cast iron pipes. I also saw crews performing additional processing on materials, such as electricians bending conduits with the "conduit bender," fire sprinkler fitters and steam fitters adding groove into the pipes with the "groover machine," or sheet metal workers applying caulk on ducts to prevent leakages. Lastly, different types of lines require different assembly techniques. I observed up guys on lifts using couplings to mechanically connect black steel or cast iron pipes, soldering or brazing copper pipes, or gluing PVC pipes. Ducts are the most laborious to assemble due to their sizes. Sheet metal workers use vice grips to temporarily attach two pieces up above and stick drift pins to keep them aligned before screwing them together on four corners with the impact gun. Unlike other MEP workers, they sometimes assemble two to three pieces on the ground for convenience and hang them simultaneously.

Salient qualities: Watching the MEP crews install lines for hours revealed that the task has all the qualities I described in the cleaning-housekeeping, fireproofing, rebar tying, and carpentry formwork tasks. Workers carrying tall, heavy materials, working on the scissor lift, and crawling around pipes or ducts to reach tight spots made me appraise this task as one of the most physical. Fire Sprinkler Fitter 1 concurred, saying, "you better be ready to realize that you're going to be paying later in life" to stress this perhaps unavoidable attribute of the task. However, identifying the task solely with this attribute would be reductive. Many of the workers' physical actions are coupled heavily with intellectual labor built on their knowledge of hardware, tools, techniques, and technical specifications, which I also learned about to a degree from them. The broader extent of this knowledge distinguishes MEP ceiling line installations from other tasks. I showed that the coupling of physical and intellectual labor in plumbers' layout work, in which reading the blueprints, committing lines and dimensions to visual memory, and taking measurements are combined with walking,

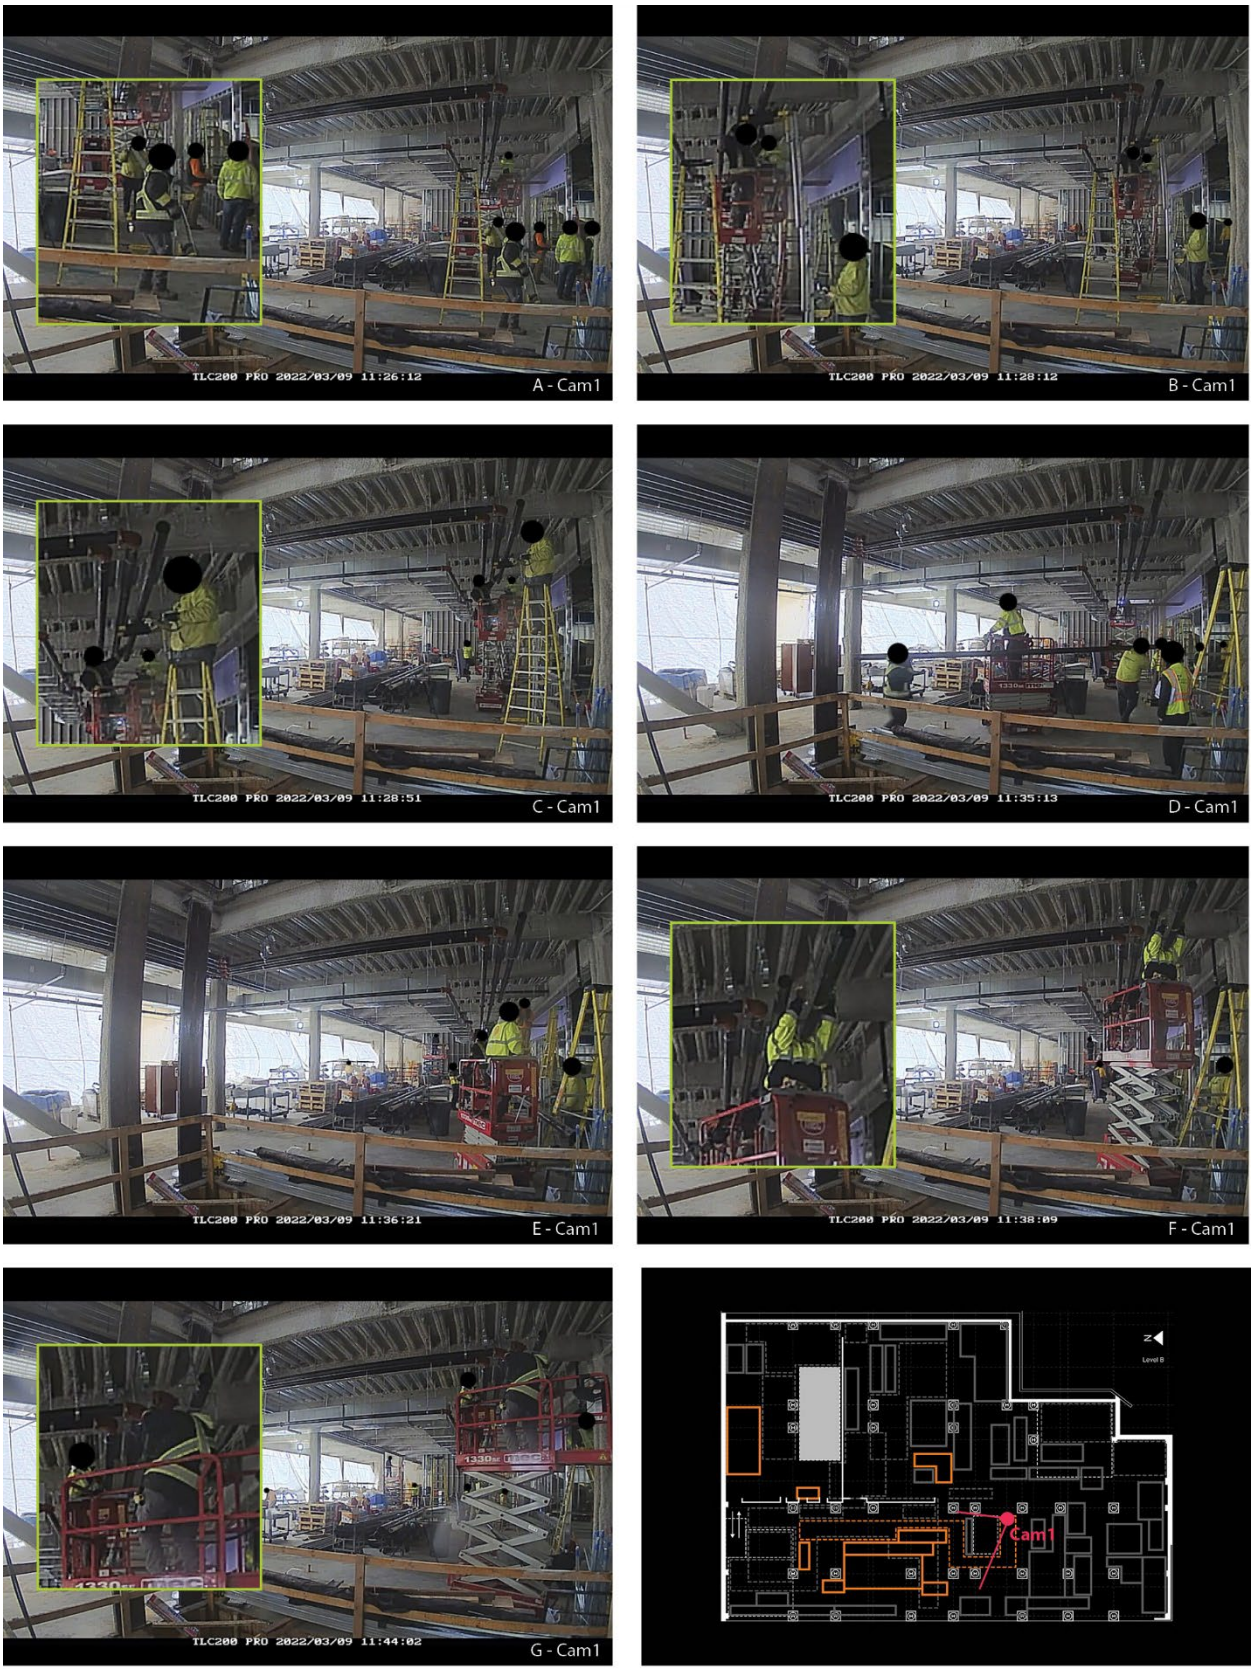


Figure 2. Key moments of Steam Fitters 1, 2, and 4's workflow on Level B

sweeping the floor, fixing the tape measure, bending down, kneeling, and marking. It is essentially a highly physical version of drafting, and the canvas is the concrete floor. Ground guys also perform a similar combination while cutting pipes and conduits; for this reason, Fire Sprinkler Fitter 1 makes his apprentices act as the ground guy first so they "look at the [blueprint], pay attention to what's going on [on the intellectual side of fitting]."

As I showed in plumbers' hanger installation and steam fitters' line fitting workflows, materials are not easy to quickly maneuver up high, the installation of each unit is done in steps, and the speed of installation is thus relatively slower than a fast-paced installation task like rebar tying. Steps, such as finding the proper elevation, fastening a rod to a strut, removing a hanger from a rod, or connecting two pieces, are all scripted and repetitive by definition. However, the set of actions performed to accomplish them is sprinkled with improvisation every time. For instance, Plumber 7 using the level to check elevations and sheet metal workers choosing to connect two to three ducts on the floor as opposed to installing one by one are improvisational moves. The expanded time span between each unit, along with improvisation, makes hanger installation and fitting appear less mechanical. In all layout, hanger installation, and fitting activities, workers are mobile, moving continuously on the ground or on lifts, but their moves are enacted mainly in a small area or in a more or less stationary position. They do not walk extensively, or cover longer distances than tasks like cleaning-housekeeping and fireproofing. They do not rapidly change location. Each of these activities involves finer, dexterous moves, such as measuring, drafting, marking, cutting, leveling, fastening, soldering, brazing, and caulking, as well as working with a rich variety of materials and tools.

Similar to fireproofing, rebar tying, and formwork, ceiling line installation is also highly collaborative, but it involves larger crews. Collaboration starts between MEP workers and fabrication shops through a "bill of material," as Plumber 4 mentioned, and fabrication shops use this document to prepare and deliver the footage for each floor. The collaboration continues between the foremen and the Superintendent doing overall planning, including the order and scheduling of how MEP workers will proceed and, accordingly, the time of delivery and the spatial arrangement of delivered batches. This requires considerable organization, in which the schedule outlining the trade, ceiling portion, and time determines the delivery dates, and the dates and the size of incoming materials guide the spatial arrangement. Ducts, main chilled water, hot water, steam, and fire sprinkler pipes cover large portions of the floor, constraining navigation. These materials are therefore distributed into areas with less MEP work but as close as possible to where they will be installed. Collaboration further occurs on the floor between and within the trades. Among the trades, coordination mainly concerns conflicts. A conflict may arise between lines—for instance, a previously installed chilled water pipe blocking the place for a duct—or between people—a steam fitter working on a pipe blocking a sheet metal worker's way. Within each trade, collaboration is about daily planning, in particular, organizing the area of operation, ordering the steps of the task in a way that maintains efficiency and minimizes conflicts, and sharing roles. Delays and improvisational changes in every aspect of planning are expected and accepted.

In many installations I observed, the roles were shared within two-to-three-people units: a ground guy reading the blueprint, doing the layout, preparing the hangers, pipes, or ducts, and passing them to one or a pair of up guys drilling and installing them on a scissor lift or ladder. Two ground-up pairs often share the daily workload in different areas. While this is typical, the unit size, roles, and choice of lift versus ladder vary based on preference as well as material sizes. For instance, instead of handling the drilling and hanger installation concurrently, Plumbers 1 and 7 preferred to do them separately and sequentially. Because hangers are light enough for one person to maneuver, Plumber 7 managed both the preparation and the installation. Similarly, PVC pipes are light enough

to be handled on ladders and installed in relatively shorter segments; consequently, I occasionally saw one plumber climbing up and down on a ladder to take care of both the preparation and the installation. Conduits are also light enough to be handled on ladders, so I observed cases in which an electrician performed it all, two electricians divided preparation and installation, or one electrician climbed up and down on a ladder to perform both while the other carried on with the installation. I saw fire sprinkler fitters doing their branch pipes in these ways too. In contrast, chilled water, hot water, and steam pipes are thick and tall. Therefore, steamfitters mostly have to use duct jacks and lifts to install them, as I described. Units consist of at least two workers on and off lifts doing both the carrying and installation or, at times, three workers sharing these roles, as I showed with Steam Fitters 1, 2, and 4. Fire sprinkler fitters for the main pipes and sheet metal crews for big ducts also work in units of this size.

Lastly, unit sizes lead to different degrees of mobility and communication. When one dedicated ground guy supports one up guy, the ground guy is more mobile, moving between the blueprint and the installation spot, bending down or on their knees for layout, or carrying, cutting, and delivering pipes. The up guy is active but stationary on a lift, fitting pipes. Communication between the workers in this case is verbal, relatively frequent, and includes the ground guy guiding the up guy about the elevations or, similar to carpentry formwork, the up guy asking for hangers, pipes, or assembly hardware. When one worker performs both the ground and up roles, many rounds of climbing up and down are inevitable. The worker is highly mobile, traversing, carrying, cutting, and fitting without any communication. Rounds of climbing up and down are also typical when two workers both take on ground and up roles, like Steam Fitters 1 and 2 did as part of the fitting I delineated. Again, the workers are mobile, performing all the preparation and installation activities together. Like in ironworkers' tying, communication between these workers is verbal and visual, reading through each other's movements. Carried out relatively frequently, it is geared towards the workers becoming one body on opposite ends of a pipe to safely lift, carry, and hang it. A dedicated up guy (Steam Fitter 4), if added to the unit, is supported by the other two. The up guy in this case is stationary on a lift, less mobile, but actively fitting. When the two workers take on ground and up roles sequentially, as was the case between Plumbers 1 and 7, they are also highly mobile, doing many rounds of up and down, but the verbal communication between them is minimal, occurring only during handovers. In all cases, familiarity or a similar level of experience reduces verbal communication.

Framing

> *he [co-worker] doesn't have to [move as much], I mean, you're on a three feet by four feet [scissor lift], and then that's all you have to go. For me, I have to run back and forth. For the bottom, I have to run back and forth. You do a lot of running. (Framer 1)*

Alongside MEP materials congesting the floor and populating the ceiling over the previous couple of weeks, another change in the scene was also difficult to miss on Level B. The floor started to be divided by unfinished interior partitions as a result of the framing and drywall installation work. Assembled by horizontal and vertical steel components called "tracks" and "studs," with fire-rated drywall panels hung on them, these partitions first separated the future mechanical room on the northeast corner from the rest of the floor. Then, a long segment running north to south delimited what would later be a large lab space on the eastern half. Across from this was a short segment running east to west and defining the entryway to a future double-loaded corridor with offices and meeting rooms on both sides. The day I planned to take a close look at framing, the two workers, Framers 1 and 2, whom I later learned were carpenters, worked on two different partitions in different areas. In the second area, they were "dropping the walls"—as it is called in construction—

that would enclose the meeting rooms on one side of this corridor (Figure 3). When I checked the workers' blueprint, I saw that this wall in the making would follow along a particular column row, which helped me project the workers' trajectory. Accordingly, I positioned a camera further down the row and began to follow their work closely.

The construction of the partition started with the layout work for its tracks on the floor, called the "bottom tracks." The two workers checked the blueprint spread on a wheeled table, located a couple of steps away from the installation area, and briefly discussed a plan. Framer 1 returned to the installation area with a broom and swept the lane on the floor where they would lay the track. He then walked around a row of assorted material stacks, which formed a continuous barrier between this lane and the western edge of the floor. Behind the stack, he tried to find a gap to extend his tape measure back to the lane. Once he did, he fixed the tape measure at a previously installed track on the western edge. This was the workers' "control line," as he told me later in an interview, that is, a reference for marking the layout for the new track. With Framer 1's end of the tape on this line, Framer 2 pulled the other end. Once he reached the point where the new track's edge would sit, he squatted down and marked it on the concrete deck with a pencil (Figure 3-A). Subsequently, Framer 1 walked a few steps down on the control line. Then, the two workers took another measurement in the same way and marked the second point. After a break caused by a conflict with a sheet metal crew in another area, they marked the third point. Once the point marks were set, Framer 2 took out a chalk box from his tool belt harness, pulled the end of a string out of the box, and handed it to Framer 1. As Framer 1 walked back to the first point mark, he positioned himself on the third one. Both workers squatted down, stretched the chalk-painted string, and touched it to the floor (Figure 3-B). The string marked the alignment line for the tracks.

According to this line, Framer 2 placed a short track previously cut down to size. Tracks usually come in ten-foot standard length and, depending on the order, in widths of either three inches and five-eighths or six inches. This piece was sized slightly longer than the width of the column it would lay next to. The piece also had a pocket cut out so it would fit around the column like a puzzle piece. To place it carefully, Framer 2 was close to the floor with one knee down. After placing the piece, he put down the barrel of a tool called the "power-actuated fastener" on it, and pulled its trigger (Figure 3-C). A nail shot from the tool and fixed the track on the concrete deck. This was a short piece, so he only nailed it on a few spots in a continuous zigzag sequence. Simultaneously, Framer 1 was cutting another short piece that would be placed next to this one. A rolling scaffold, which Framer 1 called a "moving workbench," and, later, the top surface of a gang box close to the blueprint table were his work stalls. He stabilized a standard track on the scaffold and used a power tool called the "circular saw" to trim the portion that stuck out. Additionally, he used his snips to cut out a circular hole for a pipe head. "The situation dictates what you're making," said Framer 1 when he described how they decide on the track sizes. While the workers refer to the blueprint for the track's overall length, they determine the pieces' lengths and cut-out sizes on the fly based on the sizes of obstacles. As soon as Framer 1 positioned the second small piece, Framer 2 followed with a longer one since the lane seemed uninterrupted until the next column. Walking bent down, he serially shot a few nails to fix the short track and more nails to fix the long one (Figure 3-D). In a few more rounds between the work stall and the lane performed by Framer 1, the two workers crossed another column and a few pipe heads with three short tracks. With a long piece as the final bit of the day, the bottom track spanned one-and-a-half columns.

Ideally, the installation of the ceiling track of a partition immediately follows that of its bottom tracks. However, due to ceiling tasks being conducted concurrently on Level B, the ceiling was not always open to operate on. There were instances where bottom tracks could only be installed after ceiling

tracks or where the bottom and ceiling tracks of a partition were installed at different times. The ceiling tracks of the partition that I began to observe being dropped were installed in the very early hours of the following day. As a result, I only had a chance to hear about it during my interview with Framer 1. Ceiling tracks can be painful to lay out and install depending on the MEP ceiling lines, as discussed previously. This installation also had to be performed in a tight space between a beam, a duct, and fire sprinkler pipes. Framer 1 also noted that because ceilings often feature more obstacles, such as beams or MEP lines, they require the use of short tracks and cutouts. Short tracks are also convenient to transfer from the ground to a lift top. However, they require an extra step to be attached, that is, a pair of tracks must be slid into each other for two inches and screwed. Luckily, in this case, this particular layout required only a few pieces to traverse four beams, so Framer 1 did not have to cut much. Other than a couple of plumbing pipes, he also did not have to make any cutouts. Between each beam, Framer 2, as the up guy on the lift, took measurements along a projected laser line and communicated them to Framer 1, who would do the cutting, as the ground guy. He then handed the cut pieces to Framer 2, and Framer 2 nailed them to the corrugated deck above. Whenever they arrived at a beam, Framer 1 repositioned the laser line, Framer 2 advanced with the lift, and the workers repeated their respective roles.

Once the bottom and ceiling tracks were fitted, the final step was to install the studs between them, which I followed in person. Framer 2 was again the up guy on a lift, his tool harness hanging on his side. Framer 1 was the ground guy, making rounds between a stack of studs close to one end of the track and the lift at the opposite end. The workflow began with Framer 1 marking the first few stud points with sixteen-inch intervals, on his knees. Drywall panels, which would be hung on studs, come with screw marks in rows every sixteen, twenty-four, or thirty-two inches. Studs every sixteen inches would ensure that alignment with those rows. Continuing on his knees, Framer 1 set up a point laser and transferred the first stud point to the ceiling for Framer 2's reference. He then stood up and walked to the stack, picked up one of the sixteen-foot studs, and carried it back to the point of installation. Standing in a tight spot between the lift and the bottom track, he inserted the bar's lower tip into the track, and by pivoting on that point, he lifted its entire length. In coordination with Framer 2, he positioned it vertically, straight. The two workers evaluated the bar's fit between the bottom and ceiling tracks and if any trimming was needed. Trimming is sometimes necessary because, as Framer 1 explained, "concrete is never absolutely leveled," and the workers used "deflection tracks" on the ceiling. These deeper tracks allow for a small space above studs, so the deflection of the upper deck due to weight will not flex them. The particular stud that the workers placed fit well. Thus, without affecting its positioning, Framer 2 pulled his impact gun from his harness and carefully screwed the stud into the track.

The bottom of the stud was right next to the foot of the lift and, therefore, difficult to reach with an impact gun. Further, the lift would not need to be repositioned until after a few more studs. This was why Framer 1 left the stud fixed from above only and moved on to the next point, as I learned from him on the spot. He shifted the laser to project the point, brought a new stud, and Framer 2 screwed it. Depending on a stud's positioning relative to the lift, I would see Framer 1 on either the lift's left or right side and sometimes facing it from the other side of the track (Figure 3-E). When Framer 2 screwed the fifth stud, he switched to the right side and began marking the next stud points. He was again kneeling, using the tape measure and pencil he pulled out of his tool harness. His snips, hammer, knife, and backup screws were the other tools he kept in there. After the fifth stud, Framer 2 advanced with the lift, which freed up room for Framer 1 to comfortably reach the bottom track. Framer 1 went back, double-checked whether the bottom of the studs were on the points, and screwed them to the track. He had to do this on his knees and holding his impact gun parallel to the ground. After these screws, the workers started the next set. Again, for each stud, Framer 2

repositioned the laser, walked to the stack, and brought a new stud to install. This time, he had the reach to screw the first four right away. As Framer 2 worked on the fourth, he went ahead and started the next set by screwing the first one at its bottom. Once Framer 2 joined, they installed two more studs and another one on the column corner (Figure 3-F). Framer 2 moved the lift, and Framer 1 double-checked their spacing and screwed their bottoms (Figure 3-G). Two-thirds of the frame was completed. By continuing this sequence for another hour, the two workers followed with the rest.

Salient qualities: All these activities represented a mix of qualities similar to MEP ceiling installations. Not unlike MEP layout work, framers' drafting of the lines on the blueprint on concrete is also physical and intellectual labor at the same time. Reading the blueprints, keeping the lines and dimensions in visual memory, taking measurements, and walking, positioning the tape measure, squatting, crawling, and marking all co-exist in the same activity. Broadly similar to cutting pipes and conduits, Framer 1's cutting of tracks, which entails referencing dimensions, measuring, marking, carrying, and trimming with the circular saw, also demonstrates both physical and intellectual aspects. Track and stud installation is where workers' movements become more repetitive. "Machine-like" was my first impression when I saw Framer 2 fastening the bottom tracks. Shooting nails serially is not only mechanical but also physically demanding because the worker has to put one or both knees down to be close to the target. Fixing ceiling tracks is even heavier work given that the workers deal with more pieces and wrestle with MEP lines for room. So is stud installation, when Framer 1 is standing up and kneeling on the ground and Framer 2 does the wrestling above. However, once again, being physically capable alone is not sufficient. Dexterity is needed to handle the layout, cutting, and installation swiftly. While measuring, trimming, nailing, and screwing during these activities, the framers show thoroughness and deftness with their tools.

Framing also stands out as highly collaborative, like all other multi-people tasks I covered. Similar to fireproofing and MEP ceiling installation, the task creates constraints on the floor, so collaboration begins between framers and other trades. Specifically, framers' track and stud stacks and each partition they add on the floor make it harder for other workers to navigate the space. Consequently, the overall planning of framing, in particular where framers will place the stacks, in which order they will proceed, and the scheduling of these activities are crucial. All of these points are determined via coordination with the Superintendent and other trades during scheduling meetings. The ease of framers' work depends on this coordination and the other trades they work along. "The building dictates [the work]," said Framer 1, highlighting the challenge of working in parallel with MEP workers on Level B's ceiling, as opposed to some other sites where all ceiling tracks are conveniently installed before MEP lines. In addition to overall planning, the framers' daily planning, which mainly focuses on organizing the area and the steps of the operation, relates to their location and how other trades' tasks are positioned around it. Keeping the rolling scaffold and a temporary station of tracks and studs closer while ensuring it is not disruptive to others is an important part of their planning. Not being able to operate the lift in one area will force them to do layout or bottom tracks first instead of ceiling tracks. Because the organization of the area and steps are contingent on others, improvisation is quite common.

Like in rebar tying and MEP ceiling installation, the size of the framing materials requires at least two workers to conduct the steps. Therefore, sharing roles in each step is another part of the framers' daily planning. The difference with these tasks is that Framers 1 and 2's roles are equal in some steps (like those of ironworkers) and complementary in others (like MEP ground/up pairs). Each role demands a different degree of mobility from each worker. For instance, Framers 1 and 2 take on equal roles during layout. The workers are equally mobile, performing almost the same activities, such as walking between the blueprint and the spot, taking measurements, and marking. They are



Figure 3. Key moments of Framers 1 and 2's workflow on Level B

often on their knees on the ground. They also take on almost equal roles when installing bottom tracks. Framer 1 is slightly more mobile, walking between the rolling scaffold and the spot, cutting tracks, and delivering. Then, the framers move along the track and fix it together, which is reminiscent of rebar tying. Fixing also requires being close to the ground. When doing ceiling tracks, Framer 1 becomes the supporting worker to Framer 2, who performs the installation as the lead worker on the lift. Highly mobile between the rolling scaffold, track stack, and installation spot, he guides Framer 2 in layout and receives dimensions from him to cut tracks and deliver them. Framer 2 is stationary on the lift, wrestling with MEP lines and carrying out the installation. Lastly, the roles are again equal when installing studs. Just as a horizontally positioned rebar connects the ironworkers in rebar tying, a vertically positioned stud connects the framers. Framer 1 is on the ground, more mobile and delivering the studs, while Framer 2 is again stationary on the lift. On the two ends of each stud, they are in tight coordination for positioning and fixing.

Drywall Installation

> *Topping out is a pain… I mean, you've seen, sometimes you get ten different cuts in one little four-foot section… it takes way longer. It's definitely more of a pain. (Drywall Worker 1)*

Once frames had been built in various parts of Level B, the drywall installation crew came in to hang drywall panels on these track and stud assemblages. After observing the framers building the assemblies, it was time to watch the drywall crew, Drywall Workers 1 and 2, turning these into full wall silhouettes. The two crews worked concurrently on Level B, with the drywall crew taking on what the framers finished. In earlier weeks, they worked on drywalling the frames enclosing the future mechanical room in the northeast and the future lab space in the eastern part of the floor. Later, they gradually advanced to the new frames erected in the south. One of these frames was a relatively short L-shaped segment separating the elevator core from the lab space. While Drywall Workers 1 and 2 hung panels on this frame, I set up the camera and started to observe their actions (Figure 4).

The panels the workers were to hang that day had been transferred to Level B a few days earlier. A strapped bundle of these hung from a telehandler on the ground floor and was gradually lowered to the drone test arena when I attended to watch this heavy, physical operation. Ten steps away, with an extension of the strap in hand, Drywall Worker 1 was stabilizing the bundle by standing firmly with his upper body leaned back. Framers 1 and 2 stood to his left and right, ready to slide a skateboard-like carrier, called the "skate," under the descending bundle. Drywall Worker 2 was closest to where the bundle would land, waiting. The parties synchronously approached the area as the bundle got close to touching the floor. While the drywall workers kept the load stable on opposite longitudinal sides, the framers carefully positioned the skates. The four workers then placed forty four-by-eight-foot- and five-eighths-of-an-inch-thick panels, also known as "sheetrock" or "purple board" based on their color, on these carriers (Figure 4-A). The workers untied the straps and signaled to the worker above for a second bundle. As this second one was about to stack up on the first, Framer 1 and Drywall Worker 1 picked up the wood beams on the bundle and placed it on top of the first so the two bundles would not trap the straps. The four workers made sure that the second bundle sat seamlessly on the first one and untied it (Figure 4-B). They pushed the two bundles on wheels to where they stored them in the southern half of the floor. Using a pallet jack, they lifted the bundles and transferred them on a pallet. As Drywall Worker 1 mentioned later, this location was strategic to "avoid putting it in people's way"—specifically, it was distant from MEP work and materials. This was followed by two more rounds and four more bundles.

The L-shaped segment the workers were working on required fitting custom-shaped panels to beams and assorted MEP lines in addition to full-length vertical panels at the bottom. On the day I observed the workers, they were drywalling the top, or "topping out." Like many other activities, topping out involves ground-up coordination. I previously observed this particular duo taking turns between ground and up roles, but that day, Drywall Worker 1 was on the ground responsible for shaping the panels and Drywall Worker 2 was on a lift to install them. The workers started their afternoon session by going to the stack set on the day of panel transfer and returning with two fresh panels on a specialized carrier called the "A-frame truck" (Figure 4-C). This carrier has an A-shaped neck, a backrest that holds panels horizontally standing. It works like a "mobile workstation," as Drywall Worker 1 sees it, on which he works on panels without relocating. The workflow started with Drywall Worker 2 hopping on the lift and horizontally positioning a short stud on the frame at a specific height. After attaching this additional stud to the frame with clamps, he started to take measurements along and from this baseline elevation with his tape measure (Figure 4-D). The vertical measurement was up to a pair of plumbing pipes that would have to run through the drywall. This was not merely a measurement activity but also a composition of parts that would cover an approximately four-by-eight-foot-long section and allow room for the plumbing pipes. Not too loud but in an audible voice, Drywall Worker 2 shouted the numbers indicating the dimensions of the parts he needed to Drywall Worker 1, waiting a few steps away at his mobile workstation. I realized that an interesting verbal, coordinative activity that was key to their workflow had started.

Based on the numbers, Drywall Worker 1 started to draft the parts on the panel with a T-square and pencil (Figure 4-E). From a line that he scored across the middle with his knife, he snapped the panel into two even sections and trimmed roughly one-tenth of one of these from the top. Drywall Worker 2's next announcement concerned the pipes' sizes. Accordingly, Worker 1 drew their profile with a compass based on the elevation of the pipes he had marked on the panel. He then used a spiral power saw called the "cut-out tool" and punctured what he called a "circle cut" for the pipes. His last step was slicing the panel so that the edge of each piece had half-circle cuts. Because the installed pipes could not be inserted into these closed circles, this step would allow the workers to fit the half circles around the pipes. The first piece was passed to Drywall Worker 2, and he placed it on the short stud (Figure 4-F). That was when I realized that he used this stud to set the benchmark elevation. The stud also freed his hands, so he used his left hand to align the panel with the studs behind and his right hand to hold the gun. Guided by the screw marks on the panel, he shot a series of nails (Figure 4-G). He elevated the lift to shoot the last few. The second piece would be installed higher than the previous, so Drywall Worker 1 cut it into two smaller pieces for Worker 2's convenience. He also profiled the second piece to fit the corrugated deck at the frame's topmost point. The two pieces were handed consecutively to Drywall Worker 2, and he nailed them. Because this was a double-layer wall, the workers covered the same section one more time in a slightly different composition, with more pieces. Before wrapping up the day, Drywall Worker 2 applied fire- and soundproofing chemicals around where the pieces met the corrugated deck and pipes with a pump-like tool called the "sausage gun." The workers left, leaving the remainder of the segment for the next day.

Salient qualities: This workflow nicely showed the mixed qualities of drywall installation, similar to MEP ceiling installation and framing. From transfer to installation, dealing with purple boards makes this task one of the most physically heavy alongside MEP ceiling installation. As Drywall Worker 1 pointed out, topping out, both in terms of cutting and installing, is heavier than doing bottoms, which I observed workers perform a couple of days later. Even though they install larger and sometimes uncut standard-size panels, it is "a lot less work," as he said, because it is performed on the ground with less cutting, fewer cut-outs, and no caulking. In addition to physical movements, the drywall

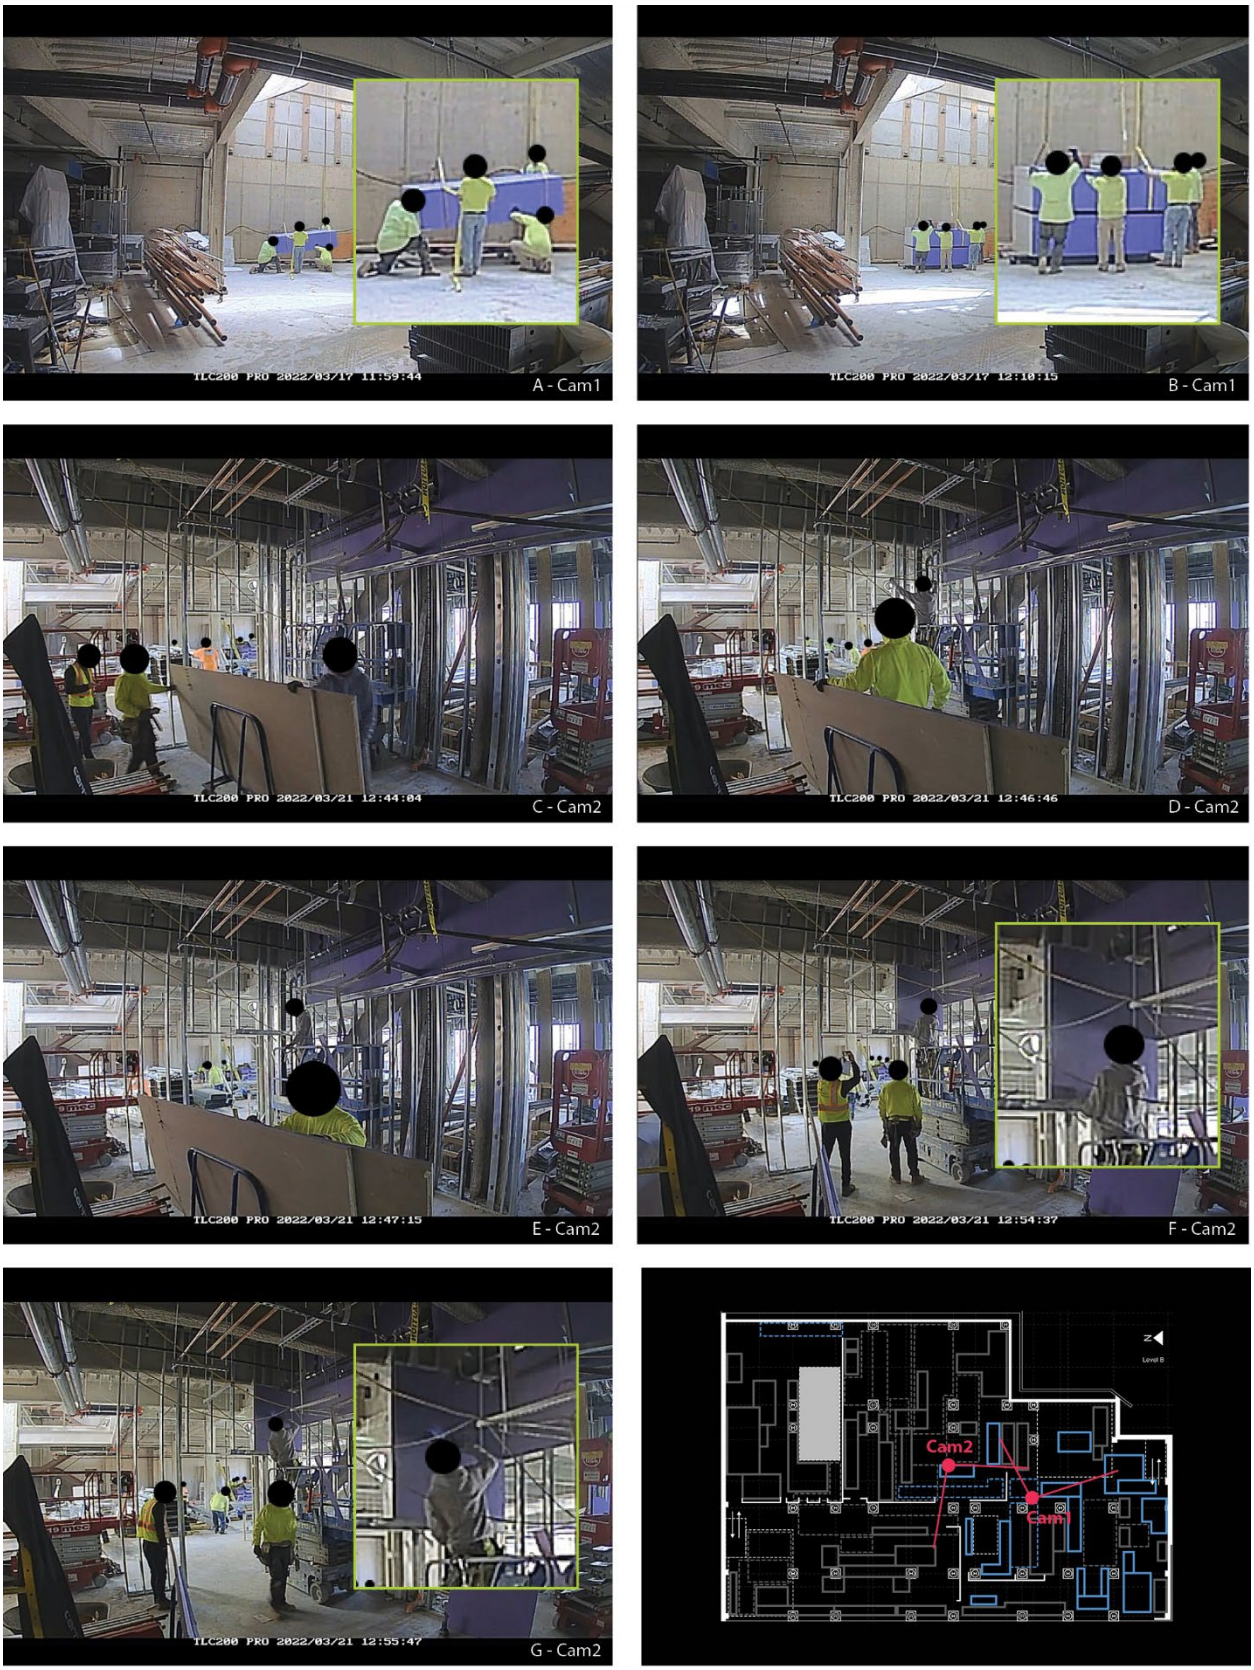


Figure 4. Key moments of Drywall Workers 1 and 2's workflow on Level B

workers' task also involves a great amount of simultaneous intellectual labor, similar to what I described in MEP ceiling installation. Setting the right elevation, measuring, and designing configurations mentally, as I watched Drywall Worker 2 do, involves spatial thinking not too different from what architects regularly perform on drawings. Similarly, Drywall Worker 1's remembering the given configuration for a four-foot section and drafting it accordingly, doing circular cuts, and producing the final pieces all involve mental agility. Usually, there is more than one possible configuration for a four-foot section, so the workers' decisions in this regard are mostly improvisational, based on experience and a simple principle: "the bigger the sheet… the faster we'll go." Other than transferring and handling the panels, almost every activity that the workers perform requires dexterity. Drywall Worker 1 displayed it as he positioned, screwed, and caulked the panels. Drywall Worker 2 showed how skillful his hands are when drafting, cutting, and making circle cuts.

Like previous tasks, drywall installation is collaborative. The task demands space for storing panels and operating the lift; therefore, planning where panels will go and where workers will proceed each week is part of the trades' coordination in scheduling meetings. Framers' progress and the MEP workers' positioning condition the drywall workers' schedule and positioning on the floor. Similar to other tasks that deal with large materials, drywall installation demands careful daily planning of the area it occupies. This includes Drywall Worker 1 setting up the A-frame station close to his co-worker so he can hear him but away from others. The workers also share roles during daily planning. Their roles and mobility change depending on whether they work on tops or bottoms. When they do tops, one is on the ground and the other on the lift, performing complementary roles. In this case, the ground guy walks more between the A-frame station and lift, delivering pieces, while the up guy installs them with his arms active but his body mainly in a stationary stance. Because tops consist of more custom pieces, the workers' progress along the frame is slower. For the same reason, the amount of time needed to complete each four-foot section varies, which makes installing tops look less mechanical. When the workers do bottoms, both become installers. One walks between the station and frame, cuts panels if needed, and positions them, while the other screws the panels. Less cutting means faster progress, and installing uncut or similarly sized panels at almost even intervals looks more mechanical. Lastly, a significant instrument is workers' verbal communication. As Drywall Worker 1 explained, the skillful way is to "know exactly where he [Drywall Worker 2] is on the sheet just by what he says," suggesting that the workers' minds and actions are perfectly synchronized. The number of cuts determines the amount of talking. "It's a lot less talking," he said, when there are fewer cuts.

Appendix D – Garbage Delivery and Concurrent Cleaning Use Cases (from Chapter 4)

The use case I saw in cleaning-housekeeping was garbage delivery to collection spots. As I described in the previous chapter, the cleaning-housekeeping worker traverses parts of or, at times, the entire floor, hand-picking garbage, sorting reusable materials, and transferring these with the pallet jack to designated spots to be dumped or stored for later use. The concept system I thought of to support this workflow would accompany the worker collecting garbage and, once loaded with garbage, autonomously deliver it to a designated collection spot in response to a prompt. This use case could nicely embody the concept of robot accompaniment. It seemed credible in the sense that a robot would have a supportive role and perform the physically heavy part of the task, while the intellectual part would remain with the worker. The delivery function was achievable with a navigation-only system, and navigation was what the group's technical efforts were already channeled into. Thus, focusing on this function would keep the group focused on navigation. The system seemed technically feasible and promised to offer a good and challenging opportunity to explore RL for task support. While all these pros crossed my mind, a conversation with Researcher 2 revealed that the system might not be as feasible as it seemed. The extensive amount of robot navigation in this case would require object detection for a large number and variety of obstacles. The lack of training datasets including construction-specific objects was a barrier. The other critical point was the size and weight of the materials that Cleaning-Housekeeping Worker calls "power garbage" and deals with. This would be an issue considering the size and payload capacity of the robots we could possibly work with.

The use case in fireproofing that I identified on the site, which was also brought up by Researcher 2 as I described earlier in this chapter, was concurrent cleaning. Scraping along lanes as Fireproofing 2 performed seemed well-suited to a robot. I imagined a lane-cleaning system that would be prompted to go to a starting point and traverse to a designated point by following along and scraping lanes. Similar to garbage delivery, this use case would also nicely demonstrate robot accompaniment. The system would become part of the duo and take over the cleaning of large areas, allowing the supporting worker to focus on gun switch and hose management duties. Like the garbage delivery system, this system would also be navigation only, which would keep the efforts focused on navigation and RL. However, during a conversation with Researcher 2, we noted two technical challenges. Having seen the fireproofing debris firsthand, I pointed out the practical challenge of scraping this thick cementitious substance with a machine. We thought that this would add equipping a commercial-grade robot with some form of mechanical scraping component as a side project. Moreover, because scraping should follow quickly after spraying, as I learned from Fireproofing 3, operating a robot closer to a scissor lift and hose would be another practical challenge. This curled, approximately two-and-a-half-inch-thick dynamic obstacle could likely cause issues in navigation. Because of these challenges, we did not explore garbage delivery and lane cleaning any further.

Appendix E – MSD-RAS Seminar Course: Key Concepts and Questions Explored

Key concepts explored in “Histories-Reconfigurations-Representations-Reflections” seminar:

ARCH 807 RAS Theory Seminar Fall 2023

Schedule

Introduction

- Week 01 I Aug 29 I Tue – Introduction + How to read, reflect, and reference
- Week 02 I Sep 04 I Mon – Labor day, no seminar

Module 1 – Histories: the intersection of design and computation

Key concepts: Taylorism, mechanization, modernism, cybernetics, automation, computer-aided design, artificial intelligence, robotics, sociotechnical imaginaries

- Week 03 I Sep 11 I Mon – Histories of the intersection
- Week 04 I Sep 18 I Mon – Technological roots
- Week 05 I Sep 25 I Mon – Reflections in the field

Module 2 – Reconfigurations: design via and around computation

Key concepts: design methods movement, design as collaborative work, sociotechnical effects, sociotechnical infrastructure, inscription, de-scription of technical objects, reconfiguration of design practices

- Week 06 I Oct 02 I Mon – Structuring design?
- Week 07 I Oct 09 I Mon – Software mediate design
- Week 08 I Oct 16 I Mon – Design via and around AI / robotic systems

Module 3 – Representations: materials in computation, materials by computation

Key concepts: digital materialities, hand-eye-tool, form, skill, craft, code, data representation, data as interface, model as analytical / generative tool, generative systems, robotic construction

- Week 09 I Oct 23 I Mon – Software as material
- Week 10 I Oct 30 I Mon – Data and model as materials
- Week 11 I Nov 06 I Mon – Materials by software, data, and model

Module 4 – Reflections: the new (and the old) vessel

Key concepts: critical reflection, CAD / AI / robotics as research site, machine learning and creativity, form making, tool making, politics of making, technological determinism, AI ethics

- Week 12 I Nov 13 I Mon – Creativity in form making, creativity in tool making
- Week 13 I Nov 20 I Mon – Politics of making
- Week 14 I Nov 27 I Mon – Technological determinism and AI ethics

---

Conclusion

- Week 15 I Dec 04 I Mon – Final presentation
- Week 15 I Dec 08 I Fri – Final paper submission

Questions articulated and explored by the students in the seminar:

- “Do off-the-shelf generative AI applications lead to new aesthetics in architecture? What are these applications’ true potentials and shortcomings for architectural design?”
- “In what ways does the early twentieth century modular architecture relate to computation? What are the implicit computational aspects embedded within some of those buildings?”
- “What are the accessibility and equity issues related to robotic fabrication in architectural education? How can the field potentially address these issues?”
- “How is collaboration in design shaped around code? What affordances, frictions, or limitations does working with code create during collaborative making?”
- “How does tradition inform contemporary design and digital fabrication of timber? What kind of knowledge that lies in wood craft is incorporated into recent projects?”
- “What are the main directions and developments in past decade’s robotic fabrication research at ETH ITA and the University of Stuttgart ICD/ITKE?”
- “What are the current directions and applications of robotics for construction in the United States’ Architecture-Engineering-Construction industry?”
- “Why is prototyping customized fabrication tools and workflows important for designers**?** What are the immediate challenges to begin?”
- “What are the drawbacks of architectural design with AI in creating experiential and emotive qualities of space? How can we safeguard and enhance these qualities?”
- “How is machine learning currently used in architectural design? What kind of data sets are employed?”