A Human-Centered, Reinforcement Learning Driven Robotic Framework For Assisting Construction Workers

SUBMITTED IN PARTIAL FULFILLMENT OF THE

REQUIREMENTS FOR

THE DEGREE OF

DOCTOR OF PHILOSOPHY

Yuning Wu

Ph.D. Candidate in Computational Design

M.S. in Machine Learning

School of Architecture, Carnegie Mellon University

Committee

Dr. Daniel Cardoso Llach

Advisor

Associate Professor, Computational Design Track Chair, School of Architecture,

Carnegie Mellon University

Dr. Jean Oh

Co-Advisor

Associate Research Professor, Robotics Institute, Carnegie Mellon University

Dr. Jieliang Luo

External Committee Member

Former Senior Principal Research Scientist, Autodesk Research AI Lab

September, 2024

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

TITLE:

A Human-Centered, Reinforcement Learning Driven Robotic Framework For Assisting Construction Workers

PRESENTED BY:

Yuning Wu

ACCEPTED BY ADVISORY COMMITTEE:

Daniel Cardoso Llach	*Advisor*	Sep 23, 2024 DATE
Jean Oh	*Co-Advisor*	Sep 24, 2024 DATE
Jieliang Luo	*External Committee Member*	Sep 26th, 2024 DATE

September, 2024

Abstract

This research explores how robots can be designed to support, rather than replace, labor-intensive workers in construction. It develops a robotic framework leveraging Reinforcement Learning (RL) algorithms, encompassing robot hardware prototyping, unstructured site perception, worker detection and tracking, hierarchical motion planning, and contextualized RL training. Realized as a "work companion robot" for carpentry workers, the framework is tested in both lab and real construction sites. Key contributions include a practical RL-driven robot for tool/material delivery and worker-robot interaction, a modular and lightweight system with tailored control and sensory packages, an comprehensive perception architecture for navigating unstructured environments, a hierarchical motion planning stack that integrates RL-based and search-based path planning, and a contextual RL finetuning pipeline for social navigation around workers. This robotic framework lays a technical foundation for future research into robotically-supported collaborative work, advancing the field of computational design, robotics, and AI/ML by integrating advanced robotics into real-world construction contexts with a focus on human-centered support.

Per aspera ad astra

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my advisory committee members. My sincere thanks go to **Daniel Cardoso Llach** for his guidance, support, and encouragement throughout this journey. Daniel has been both patient and insightful, and his feedback has been crucial to the completion of this work. Looking back over the past five years, I have learned much more from him than just academics. Daniel has taught me the importance of being collegial, keen, humble, and open-minded to opportunities and inspirations in both research and life.

I am also deeply grateful to **Jean Oh**, my co-advisor, for mentoring my journey in robotics and AI. Jean has been incredibly supportive and encouraging, providing critical advice at key moments. Her wisdom in navigating both research and future careers has offered me lifelong and profound insights.

Additionally, my heartfelt thanks to **Jieliang (Rodger) Luo**. Our collaboration began during my internship at Autodesk Research, and I am thankful for his support in many aspects of my PhD over the past years. His critical insights and feedback have been invaluable to my research and career development. I also wish to express my gratitude to the amazing colleagues at Autodesk AI Lab and Robotics Lab, particularly **Hui Li**, **Tonya Custis**, and **Sachin Chitta**. Your support for my research during that summer has

been a great inspiration for my PhD journey.

I would also like to thank colleagues and teammates in the ReAC Research Group, particularly **Jiaying (Vina) Wei** and **Emek Erdolu**. Vina participated in the research during her master's study and has been with me throughout the entire process. She has been a great research collaborator and friend, and I am grateful for her support and encouragement, especially when the robot was malfunctioning. The long hours spent building a robot together and making it operable are invaluable memories. Emek has been a inspiring colleague, and I am thankful for his insights on the ethnographic study. I would also like to extend my thanks to all colleagues who have participated in Husky's building lifecycle, including **Sam Shum**, **Meghdeep Jana**, **Tingsong Ou**, **Yuchen (Joshua) Cao**, **Haoyu Liu**, **Linxiaoyi Wan**, and others.

I would like to express my gratitude to my family for their unwavering love and support. The encouraging calls across the Pacific Ocean gave me the strength to persevere through this challenging journey. If I may call this pursuit a marathon, you are the water stations along the way.

I am also deeply grateful to my dear friends from Pittsburgh, **Xiaohan Zhou**, **Anni Li**, **Yiwen Jiang**, **Junyi Zhu**, **Xuehui Zhang**, **Yuchuan Shan**, **Yiwei Huang**, and others. Your companionship made the five years in Pittsburgh, particularly during the pandemic and a critical surgery away from my family, bearable. Whenever I reflect on my PhD journey and the city, my fondest memories include playing Catan, going out for barbecues, and engaging in film photography.

Special thanks to my friends from Los Angeles, **Yue Shen**, **Yueling Chen**, and others. Your presence, along with the California sunshine and beaches, provided the perfect respite and inspiration. I often refer to my trips to Los Angeles as solar recharging, and the late-night landings at LAX

were always moments of hugs and joy. A true friend is someone who wouldn't hesitate to pick you up at LAX in late-night traffic across town. Because of you, the city feels like a cozy, secret second home for me.

A special section of gratitude goes to my cats, **Sophie (Tangtang)** and **Mister (Huihui)**. Your presence and companionship have been a great source of comfort and joy. The daily routine of feeding, playing, and cuddling with you has become an important part of my life.

I am grateful for the support I received in the interdisciplinary research environment here at Carnegie Mellon University. I would like to acknowledge the faculty and colleagues from **CMU Manufacturing Futures Institute** at Mill-19, particularly **Rod Heiple**, **Sandra DeVincent Wolf**, and others. Working at the Mill-19 lab surrounded by robots for a whole day is one of the best memories throughout my PhD. I also wish to extend my gratitude to faculty and colleagues from the **School of Architecture, the Robotics Institute, Computational Design Laboratory, Bot Intelligence Group, National Robotics Engineering Center**, and **Machine Learning Department**. I constantly remind myself that my current academic gains are built upon your kind support. The hardworking research aura of this renowned institute has greatly honed my research skills, motivations, and methodologies. Additionally, I have also been generously supported by **Autodesk Research**, **State of Pennsylvania**, **TCS**, and **LEGO Foundation**, through scholarships and fellowships. Thank you for trusting me and providing me with the resources to pursue my research interests.

A sincere thank you to my former advisors and colleagues from Tsinghua University, particularly **Rong Rao**, **Xiang Dong**, **Rong Zhou** who saw my potential beyond traditional architectural design and encouraged me to apply my instincts in system architecture, design, and spatial creativity to

explore the interdisciplinary realms of robotics and AI. Nothing beats the joy of discovering true research passion.

This work is a culmination of all the support, encouragement, and hard work from everyone mentioned above. Thank you all.

Glossary

agent An agent in reinforcement learning is the entity that interacts with the environment. It receives observations and rewards from the environment, and selects actions to perform.. 25

deep neural network A Deep Neural Network (DNN) is an artificial neural network characterized by multiple layers between the input and output layers. Each layer consists of interconnected neurons that apply transformations to input data, enabling the network to learn complex patterns and representations [LeCun et al., 2015].. 20

Lidar Odometry and Mapping (LOAM) Lidar Odometry and Mapping (LOAM) [Zhang and Singh, 2014] involves using lidar data to estimate a robot's position and generate a map of the environment, critical for precise navigation.. 22

localization Localization is the process of estimating a robot's pose (position and orientation) in a known map of the environment. It is a key component of robot navigation.. 22

Machine Learning (ML) Machine learning is a field of artificial intelligence that focuses on developing algorithms and models that can learn from data and make predictions or decisions without being explicitly programmed.. 25

mapping Mapping is the process of creating a map of the environment using sensor data. It is a key preliminary component of robot navigation.. 22

navigation stack A navigation stack is a set of software components that work together to enable a robot to navigate in an environment. It typically includes components for localization, mapping, path planning, and obstacle avoidance.. 54

obstacle avoidance Obstacle avoidance is the task of detecting and navigating around obstacles in the environment to prevent collisions. It is an essential aspect of robot navigation.. 88

path planning Path planning is the process of finding an optimal path from a start to a goal location while avoiding obstacles in the environment. It is a key component of robot navigation.. 21

policy A policy in RL is a mapping from states to actions. It defines the agent's behavior in an environment, and can be seen as the agent's decision making strategy.. 25

Reinforcement Learning (RL) Reinforcement Learning (RL) is a type of machine learning where an agent learns decision-making by interacting with an environment. The agent receives feedback in the form of rewards, and its goal is to learn a policy that maximizes the cumulative reward over time [Sutton and Barto, 2018].. 25

reward A reward in RL is a scalar value that represents the feedback from the environment to the agent. The agent's goal is to maximize the expected cumulative reward it receives over time.. 25

simulation environment A simulation environment is a software-based environment that simulates the dynamics of a real-world environment. It is used for training and testing RL algorithms.. 28

social robot navigation Social robot navigation is the task of navigating a robot in a social environment where it interacts with humans and other robots. It requires the robot to learn social norms and conventions to navigate effectively. In this research, it is also referred to as social navigation.. 19

state A state in RL is a representation of the agent's environment at a given time, encompassing all relevant information needed for decision-making, such as positions, velocities, and sensory data.. 58

state estimation State estimation is the process of estimating the state of a system based on sensor measurements and control inputs. In the context of robot navigation, it refers to estimating the robot's pose and velocity.. 73

value A value (function) in RL is a function that estimates the expected cumulative reward an agent can achieve from a given state or state-action pair.. 27

Contents

1 Introduction **1**

1.1 Research Question and Hypothesis Outline 3

1.2 Thesis Structure 4

2 Background **7**

2.1 Trends in Construction Robotics 8

2.1.1 Material-Centric Trend 8

2.1.2 Robot-Centric Trend 11

2.1.3 Design-Centric Trend 16

2.1.4 The Emerging Human-Centric Trend 18

2.2 Developing Socially-Aware Navigation Robots 19

2.2.1 Site Perception 22

2.2.2 Detection and Tracking of People 23

2.2.3 Path Planning 24

2.3 Reinforcement Learning for Adaptive Robots 25

2.3.1 Reinforcement Learning in Construction Robotics . . 26

2.3.2 Different Types of RL Algorithms 27

2.3.3 Simulation Environment for RL 28

2.4 Context of Carpentry Formwork 30

3 Hypothesis **35**

3.1 The Robot Support Scenario 37

3.2 Specificity of site and material 40
3.2.1 Site Objects and Conditions 40
3.2.2 Carpentry Formwork Tools and Hardware 44
3.3 Specificity of Formwork Activities 48
3.3.1 Individual Work Paradigms 49
3.3.2 Collaborative Work Paradigms 51
3.4 Context-Specific Technical Challenges 53

4 Methods **56**
4.1 Robot Hardware Design 62
4.1.1 Base and Chassis 64
4.1.2 LiDARs 67
4.1.3 Cameras 69
4.1.4 Additional Structure and Design 71
4.2 Unstructured Site Mapping and Robot State Estimation . . . 73
4.2.1 Mapping 74
4.2.2 Robot State Estimation 76
4.3 Worker Detection and Tracking 78
4.3.1 Detection 79
4.3.2 Tracking 84
4.4 Hierarchical Motion Planning 85
4.5 Contextualized Fine-tuning of RL 90
4.5.1 Customized Agent Behavioural Patterns 93
4.5.2 Customized Agent Layout Configurations 94
4.5.3 Limited FoV of the Robot 97
4.5.4 Domain Randomization 97
4.5.5 A Uniform Pipeline For Incremental Model Adaptation 98

5 Results **102**
5.1 Final Robot Prototype 103

5.2 Simulation Training Results 104
5.3 Qualitative Evaluation On Site 105
5.3.1 Site Setup 105
5.3.2 Worker Interaction Scripts 106
5.3.3 Pre-Demonstration Preparation 110
5.3.4 Robot Demonstration Performance 113
5.3.5 Post-Demonstration Worker Interview 120
5.4 Quantitative Evaluation in the Lab 122
5.4.1 Lab Sandbox Setup 122
5.4.2 Evaluation Scenarios and Test Cases 123
5.4.3 Evaluation Metrics 126
5.4.4 Quantitative Evaluation Results 126

6 Reflection and Discussion 130
6.1 Locating Robotic Technology Design: The Who, The What, and The How 130
6.1.1 The Who: Identifying the Users and Their Work Activities 131
6.1.2 The What: Localizing the Specific Robotic Support Task 132
6.1.3 The How: Realizing Robotic Support in a Contextually Informed Manner 133
6.2 Architecting A Robot For Real-World Challenges 137
6.3 "Hardware is Hard" 139
6.4 Is Reinforcement Learning the Answer? 142

7 Conclusion and Future Work 146
7.1 Contributions 146
7.2 Limitations 150
7.3 Future Work 151

Worker Interview Summary **176**

Related Publications **181**

List of Figures

2.1 The textile form structures by Sean Ahlquist [Ahlquist et al., 2013, Ahlquist and Menges, 2013, Ahlquist, 2015], in which a series of computational methods are applied towards a thin and delicate material expression. 9
2.2 The distributed timber framework by ICD Stuttgart [Leder et al., 2024], an example of material-centric construction robotics. 10
2.3 Robotic drones for 3D concrete printing [Zhang et al., 2022] and bricklaying [Fang et al., 2020]. 12
2.4 Tightening timber joints with robotic arms [Apolinarska et al., 2021]. 14
2.5 Field mobility robots for autonomous excavation [Johns et al., 2023] and layout and surveying [Lau, 2020]. 15
2.6 Design-centric research in construction robotics, featuring Gilles Retsin's installations 17
2.7 Examples of human-centric research in enhancing physical work capabilities, including exoskeletons [Nuckols et al., 2021] (left) and haptic/UI robotic systems [Brosque et al., 2021] (right). 19
2.8 Carpentry formwork activities observed on-site. Image credit: ReAC Research Group, 2022, diagram and photo by Emek Erdolu Erdolu [2024]. 33

3.1 Robot support scenarios and functions: sending, summoning, and accompanying [Wu et al., 2024b]. 38
3.2 Four key categories of site objects, (1) low, temporary, shape-varying objects, (2) low to medium height, bulky objects, (3) high, fixed installations, and (4) high, thin, temporary objects. 42
3.3 Case of open twin ladders: (A) scanning at a low height risks the robot passing underneath the ladder; (B) scanning at a high height risks the robot being tripped by the ladder; (C) shows the desirable 2D scan representation for collision-free navigation. 44
3.4 Site photos of tools and hardware in carpentry formwork. Image credit: ReAC Research Group, 2022, photos by Emek Erdolu. 45
3.5 The three individual work paradigms: (1) work and progress by the formwork wall, (2) traversing the site, and (3) installing and standing still. 51
3.6 The two collaborative work paradigms: (1) interaction and (2) separation. 53

4.1 System design of the RL-driven robotic framework [Wu et al., 2024b]. 60
4.2 Three major stages in robot hardware design and development. 63
4.3 Comparison among three robot bases. 65
4.4 "Blind cone" effect caused by the metal rims of a 3D LiDAR. Image credit: ReAC Research Group, 2022, renderings by Jiaying Wei. 68
4.5 Composite LiDAR box design (left) and implementation (right). 69
4.6 Tilted camera stand design and the duo-camera setup. 70
4.7 Retractable tool and hardware container. Image credit: ReAC Research Group, 2022, rendering by Jiaying Wei. 72

4.8 Custom-designed Bluetooth controller. 73
4.9 Three iterations of site mapping: (A) simple 2D mapping; (B) 3D mapping and point cloud reconstruction; and (C) processed 2D mapping from real-time point cloud reconstruction. Adapted from [Wu et al., 2024b]. 77
4.10 The entire worker detection and tracking process including (A) raw camera detections, (B) detection after clustering, and (C) the ultimate tracking result. 86
4.11 The dual-layered hierarchical motion planning stack. 87
4.12 The customized agent patterns. 95
4.13 The customized agent layout configurations. 97
4.14 Overall system dashboard of the robotic framework. For further clarity, a larger version of this image can be accessed online: `https://ice-5.github.io/assets/img/thesis/dashboard.png`. 101

5.1 Final robot prototype. This image was previously published in [Wu et al., 2024b] 104
5.2 Training process of the contextual fine-tuning of the RL model in the simulation environment, shown by cumulative reward (left) and collision rate (right). A similar version of this image appeared in [Wu et al., 2024b]. 105
5.3 On-site demonstration and evaluation setup. The photo was also covered in [Wu et al., 2024b]. 107
5.4 Storyboard of the three scripted scenarios. 109
5.5 Mapping of the demonstration site. Part of this image was included in [Wu et al., 2024a]. 111
5.6 Rapid installation of the third camera in the lab. 113

5.7 Robot trajectory from the time-lapse recording and photo documentation of Script 1. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a]. 116
5.8 Robot trajectory from the time-lapse recording and photo documentation of Script 2. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a]. 117
5.9 Robot trajectory from the time-lapse recording and photo documentation of Script 3. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a]. 119
5.10 Laboratory evaluation setup. This image was included in [Wu et al., 2024a]. 124
5.11 The two evaluation scenarios during quantitative evaluation: (left) scenario primarily tests the robot's navigation support capability among work zones by the target wall; (right) scenario focuses on the delivery capability between the central workbench and work zones. 124
5.12 Photos captured during the lab evaluation. Some photos were published in [Wu et al., 2024b]. 127
5.13 Trajectory comparison between non-RL (left) and RL (right) based robot navigation. This image result was published in [Wu et al., 2024b]. 128

7.1 A preview and current progress of developing a photorealistic, physics-based simulation environment for benchmarking context-awareness of robotic system in work environment. Image credit: ReAC Research Group, 2023. 154

List of Tables

3.1 Size and weight estimate of carpentry formwork tools and hardware. 46

5.1 Quantitative results in lab evaluation regarding context-specific, comfort-related metrics [Wu et al., 2024b]. 127

5.2 Quantitative results in lab evaluation. "Success": successful delivery rate. "Collision": major/minor collision rate. "Freeze": robot freeze rate. "Time": navigation time (seconds) between two benchmarking work zones [Wu et al., 2024b]. 128

Chapter 1

Introduction

"Will robots replace humans?" This is a question I frequently encounter during my morning Uber rides to the Mill-19 Robotic Research Center when people learn about my research. Beneath this question lies widespread speculation about the future of robots. With the advent of advanced AI technologies like Large Language Models (LLM), there is a prevailing ambition to develop AI-driven, autonomous robots to replace human workers, which often evokes concerns about a jobless future. In response to this question, my research investigates an alternative vision of robots and automation: a human-centered approach. By human-centered, I mean more than just a friendly stance toward existing human efforts. It critically envisions a collaborative work paradigm where human work capabilities are practically supported and enhanced by robots.

To substantiate this vision, this research targets the construction industry, one of the world's most challenging, labor-intensive, and least automated sectors. Unlike manufacturing pipelines housed in organized factories, most construction activities still occur on-site, with workers navigating bespoke building requirements in complex environments. These manual practices

often involve significant physical strain and continued fatigue. Traditional approaches to automating construction, which rely on replacing individual manual tasks with single-purpose machines, face difficulties due to the diverse, improvisational, and complex nature of daily construction work.

This research challenges this conventional view of robots as human surrogates, proposing instead a supportive integration of robotics into existing construction workflows. Using a "research by prototyping" approach, the study develops a reinforcement learning-driven robotic framework that enables robots to adaptively integrate into the intricacies of the existing workflows, offering nuanced assistance to workers rather than reconfiguring or displacing them. Focusing on carpentry formwork as a representative context, the framework is grounded in the social and material specificities of this work practice. Developed in tight coupling with my role in the Rethinking Automation in Construction (ReAC) group at Carnegie Mellon University SoA's Computational Design Laboratory, the framework aims for safe and socially aware robot mobility around workers engaged in formwork installation, helping carry and deliver essential but heavy tools and materials in close proximity. Key contributions of this research are (1) an RL-driven "work companion robot" prototype, a physical realization of the framework, for real construction environments, (2) a modular and lightweight robotic system with tailored control and sensory packages, (3) an advanced perception architecture for navigating unstructured obstacles and detecting worker activities, (4) a hierarchical motion planning stack that integrates RL-based and search-based methods, and (5) a practical pipline of contextualized RL finetuning for socially adaptive navigation around workers. The framework is evaluated both in the lab and on an actual construction site to assess its ability to provide practical, comfortable assistance, and lays a technical foundation for future investigations into the broader algorithmic, ethical, social, and applicational aspects of robotically-supported collaborative work.

An inherently interdisciplinary quest, this research intersects the study of social robot navigation, reinforcement learning, and construction workflows. It explores a human-centered perspective that previous computational design researchers have seldom emphasized in investigating how we build and create physical spaces. The research designs and develops robotic technologies drawing from the "always-already" [Heidegger, 2010] social and material characteristics of their daily work.

This research also taps into the rich body of existing and ongoing work in human-centered AI [Shneiderman, 2022, Krishna et al., 2022], Human-Robot Interaction (HRI) [Breazeal, 2004, Bethel and Murphy, 2010, Seaborn et al., 2023], and embodied AI [Savva et al., 2019, Duan et al., 2022]. It contributes to these broader research areas by showcasing sensitivity and consciousness regarding the situational context of applying these technologies. Specifically, the research demonstrates how the development of robotics and AI/ML technology can be grounded within and inspired by social and material insights drawn from an in-depth study of the fabric of actual human work.

1.1 Research Question and Hypothesis Outline

Rather than replacing humans, how can advancements in RL and robotics be leveraged to create robots that can adaptively support, accompany, and assist workers in their existing labor-intensive workflows in on-site construction contexts?

In response to this research question, this research proposes that RL-based social navigation methods can work in tandem with (1) robot system design and hardware setup, (2) unstructured site perception and robot state estimation, (3) worker detection and tracking, (4) efficient and hierarchical motion planning, and (5) contextualized fine-tuning of RL to collabora-

tively support on-site construction workers and realize navigation comfort and adaptability in effective, measurable ways. Synthesized as an RL-driven robotic framework, it should (1) be able to enable safe and fluent robot navigation on an existing complex construction site, populated with busy workers, (2) observe and adapt to workers' movements and activities such that close-range worker-robot interactions proceed with a level of social comfort, and (3) be able to offer meaningful forms of live work assistance that are helpful in a labor-intensive construction workflow.

1.2 Thesis Structure

The subsequent chapters are organized as follows. Chapter 2 reviews existing research and developments in construction robotics, socially-aware navigation robots, and reinforcement learning. It identifies three major trends in construction robotics: material-centric, design-centric, and robot-centric, while also noting the emerging human-centric research trend. The chapter then discusses the evolution and current state of socially-aware robot navigation, covering key algorithms and necessary technical components. Following this, it examines the role of reinforcement learning in advancing construction robots, social robot navigation, and computational design. The chapter concludes by situating the research within the context of carpentry formwork.

Chapter 3 Hypothesis delves into the robot support scenario, the specificity of the site and materials, and the specific activities involved in formwork installation [1]. This context setting is crucial for understanding the challenges and requirements for developing the supportive robotic frame-

[1]The analysis draws on ethnographic and qualitative studies conducted by the ReAC Research Group, incorporating on-site recordings, worker interviews, and preliminary documentation of work patterns.

work.

Chapter 4 Methods describes the comprehensive approach taken to develop the robotic framework, including the design and development of the robot hardware, methods for unstructured site mapping and robot state estimation, worker detection and tracking, hierarchical motion planning, and the contextualized fine-tuning of reinforcement learning algorithms. This chapter is essential for understanding how the research objectives are pursued.

Chapter 5 Results presents the outcomes of the research, including the final robot prototype, qualitative and quantitative evaluations conducted both on a construction site and in the lab, and additional findings that emerged from the research process.

Chapter 6 Reflection and Discussion provides critical reflections on the technology design and development process through the lenses of the "who", the "what", and the "how" in the research. It also examines the architecture and system integration of the robot, discussing the challenges encountered in both software and hardware development. Furthermore, this chapter delves into the technical effort required to adapt robot learning methods to specific contexts and evaluates whether reinforcement learning effectively addresses these challenges.

Chapter 7 Conclusion and Future Work summarizes the research findings and contributions, including robot hardware prototyping, unstructured site perception, worker detection and tracking, hierarchical motion planning, and contextualized RL training, and the implications and limitations of the work. It also outlines potential directions for future research in "robotically-supported collaborative work".

The appendix includes a summarized transcript of the worker interviews,

providing further insights and data to support the research findings.

Chapter 2

Background

Reflecting the interdisciplinary nature of this research, this chapter presents a comprehensive review of the literature on three associated domains: construction robotics, social mobile robots, and reinforcement learning. Major trends and gaps in each of these domains are discussed in relation to this research, highlighting inspiration from promising aspects while identifying unaddressed challenges. The chapter is structured as follows: Section 2.1 provides an overview of construction robotics, highlighting key trends and advancements in the field. Section 2.2 discusses the development of socially aware navigation robots, also covering supplementary technical issues on the detection and tracking of people, perception of surroundings, socially aware local path planning, and simulation environments. Section 2.3 explores the role of reinforcement learning in enabling adaptive robots, emphasizing its potential applications and challenges in the aforementioned domains. Section 2.4 provides context for carpentry formwork, outlining the workflow and explaining why it was chosen as a task suitable for robotic support. The chapter concludes with a summary of key insights and research gaps identified in the literature review.

2.1 Trends in Construction Robotics

The field of construction robotics has undergone continuous research over the past decades [Bock, 2007, Xiao et al., 2022, Attalla et al., 2023], gaining increased momentum with the recent surge in new computational methods and tools. Although there are diverse approaches, demonstrations, applications, and impacts in this field, several predominant trends are observed, especially in recent years' research. These trends critically examine essential elements in designing and creating physical spaces, including materials, robots/machines, design, and the growing focus on human factors. Respectively, the following review categorizes and articulates these trends as material-centric, robot-centric, design-centric, and the emerging human-centric trend.

2.1.1 Material-Centric Trend

Since construction fundamentally involves the handling and structuring of materials, some scholars in architecture and civil engineering adopt a material-centric perspective towards robotic applications. This approach emphasizes the characteristics of materials, viewing them "not as passive recipients of form but as active sources of inspiration and creativity" [Menges, 2015]. This perspective has spurred the use of emerging computational methods, tools, and robots for exploring the properties, expressions, and prototypes of various materials. Within this research framework, robots are primarily seen as *agents for prototyping*, valued for their ability to semi-autonomously or autonomously handle complex and precise material tasks, some exceeding traditional human craftsmanship. A notable representative of this trend is Achim Menges and his research group at ICD Stuttgart, who advocate for materiality as "an active generator of design and an adaptive agent of

architectural performance" [Menges, 2015]. They have developed various material systems, prototypes, and demonstrations using innovative robotic technologies, significantly influencing peers in this research area.

For example, in Sean Ahlquist's "textile hybrid structure" [Ahlquist et al., 2013, Ahlquist and Menges, 2013, Ahlquist, 2015], computational methods like experimental form-finding, spring-based simulation, and finite element analysis are applied to physically examine the behaviors of force-active membranes and elastic, thin rod combinations (Figure 2.1). Materials such as open-web meshes, CNC knitted textiles, and fiberglass rebar rods are used to showcase an expandable, balanced structural system.

Figure 2.1: The textile form structures by Sean Ahlquist [Ahlquist et al., 2013, Ahlquist and Menges, 2013, Ahlquist, 2015], in which a series of computational methods are applied towards a thin and delicate material expression.

Another example is the distributed robotic timber assembly system [Leder et al., 2019, 2024], where a fleet of small, mobile robots demonstrates their utility in forming a timber-beam material system (Figure 2.2). These robots are designed specifically for the size and weight of timber beams and serve as both material manipulators and parts of the timber framework joints. The robots use low-cost servo motors for basic rotation and strut-grasping actions, acting as flexible and adjustable connectors for adjacent timber beams in diverse connecting scenarios. This approach reformulates the construction of complex, large-scale timber structures.

Figure 2.2: The distributed timber framework by ICD Stuttgart [Leder et al., 2024], an example of material-centric construction robotics.

Beyond these examples, similarly intended robotic applications are observed in the experimentation and construction of various other materials, including composite shell structures created through additive manufacturing [Felbrich et al., 2018], long-span structures fabricated collectively by industrial robot arms and UAVs [Felbrich, 2017], eroded stone structures made with event-based robotic control [Micoli et al., 2018], and filament structures installed by mobile robotic fabrication systems [Yablonina et al., 2017]. Despite the diversity in robot types[1] and computational techniques[2], the envisioned roles of robots are closely tied to the material forms and features they are designed to build or produce.

The material-centric trend has indeed infused new inspiration into construction robotics and computational design, highlighting the untapped potential of materials and their physical attributes in the design and creation of the built environment. However, this trend also raises questions about the scalability, adaptability, and practicality of these material-centric systems in existing, real-world scenarios. The elegant artifacts created in these research projects, often exhibited in art biennales, museums, galleries, or as singular statues or pavilions, are still distant, if not disconnected, from actual, messy, and commonplace construction sites and practices. The gap

[1]Such as drones, robotic arms, mobile ground robots, etc.
[2]Such as additive manufacturing, robotic manipulation, robotic navigation, etc.

between lab-based material prototypes and on-site construction practices remains a challenge to be addressed, necessitating further research to bridge this divide.

2.1.2 Robot-Centric Trend

Another major trend in the robotics and construction community is robot-centric. This trend primarily focuses on replacing labor with autonomous robots, machines, and systems to enhance efficiency and productivity. Robots are thus envisioned as *surrogates for humans*, tasked with automating specific construction procedures or workflows with minimal human intervention. Current research and application efforts are concentrated on tasks that are repetitive, physically intense, but relatively less skillful, including robotic spray painting [McAlorum et al., 2021, Ercan Jenny et al., 2022], brick-laying [Bonwetsch, 2015, Fang et al., 2020], screed-leveling [Melenbrink et al., 2020, Yang et al., 2021], masonry [Goessens et al., 2018], and concrete printing [Borg Costanzi et al., 2018, Hahlbrock et al., 2022].

Among these robotic experiments, drones are often leveraged for their mobility and flexibility in navigating construction sites. For instance, there are studies leveraging drones for aerial surveying [Anwar et al., 2018], site inspection [Ashour et al., 2016], and material lifting and transportation [Choi and Kim, 2019] in construction. In recent work by researchers from UCL and UPenn, a swarm of drones is designed for collective, on-site 3D printing of concrete [Zhang et al., 2022] (A in Figure 2.3). Inspired by bees, each individual drone carries a small 3D printer, and collectively, they can coordinate and print large-scale, complex-shaped concrete structures in a distributed manner. A similar idea of a distributed drone system is also explored in the ReAC group, where a scalable fleet of drones is targeted for semi-autonomous bricklaying and spray-coating tasks, leveraging recent

multi-agent reinforcement learning (MARL) algorithms for collision avoidance and task coordination [Fang et al., 2020] (B in Figure 2.3). While these drone systems are still in the experimental stage, they showcase the potential of aerial robots in automating tasks that are traditionally labor-intensive, time-consuming, or even infeasible for human workers. However, drawing from the ReAC group's firsthand experience, the futuristic image of such multi-drone autonomy must, in fact, be sustained by an entire team of human operators, engineers, and researchers, who are responsible for the drones' supply, maintenance, and safety control. These hidden yet critical human efforts necessitate viewing the narrative of autonomous robots in these cases as a complex interplay between human and machine.

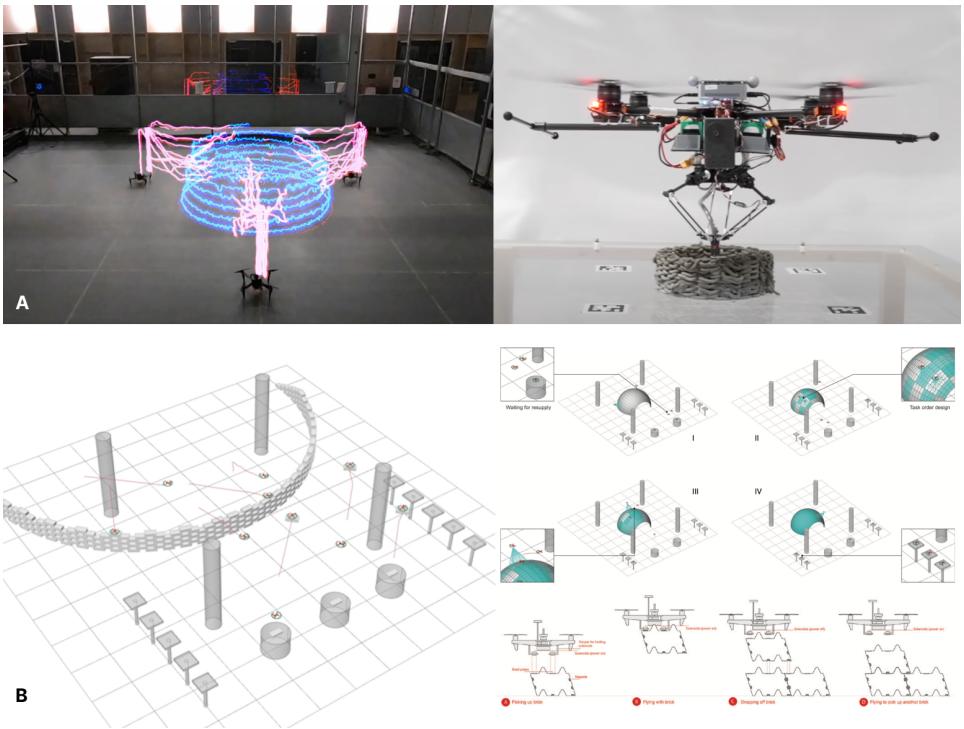

Figure 2.3: Robotic drones for 3D concrete printing [Zhang et al., 2022] and bricklaying [Fang et al., 2020].

Another significant area of research and application in the robot-centric trend involves using robotic arms for contact-rich manipulation in construc-

tion. These robots perform tasks requiring precise, dexterous, and adaptive interactions with construction tools, materials, and equipment. For example, researchers at ETH Zurich have extensively used robotic arms for timber assembly, where robots are programmed or trained to semi-autonomously perform joint assembly [Willmann et al., 2016, Thoma et al., 2019, Apolinarska et al., 2021] (Figure 2.4). This includes fundamental tasks such as picking up a beam and placing it within an etching on another beam [Apolinarska et al., 2021], or positioning it for workers to secure with nails. Another notable project explores using robotic arms for on-site plasterwork [Ercan Jenny et al., 2020], where a robotic arm is mounted on a mobile base to enable continuous plastering along extended wall surfaces. This robotic arm is equipped with a plaster nozzle and a vision system to monitor the plastering process in real-time, ensuring consistency and quality.

These examples highlight the potential of robotic arms in automating construction tasks by emulating certain human manipulation behaviors. While promising, it is important to acknowledge that the scope and subtlety of actual human assembly work cover more than a well-defined, atomic, joint insertion task. Still taking timber assembly as an example, although the robot can tighten the joint, it lacks the ability to handle nuances associated with unexpected situations, make minor corrections, contingently fortify the structure by adding a nail, adaptively trim and adjust the etches for better joint fitness, or promptly conduct quality inspections—typical practices in a human worker's daily tasks. These often neglected nuances, a sense of situational awareness, are crucial components of the rich and adaptive fabric of actual human work, but are often overshadowed by the celebration of success in one or two major robotic tasks.

Besides drones and robotic arms, other types of robots concerning field mobility are also explored in the robot-centric trend. For instance, recent

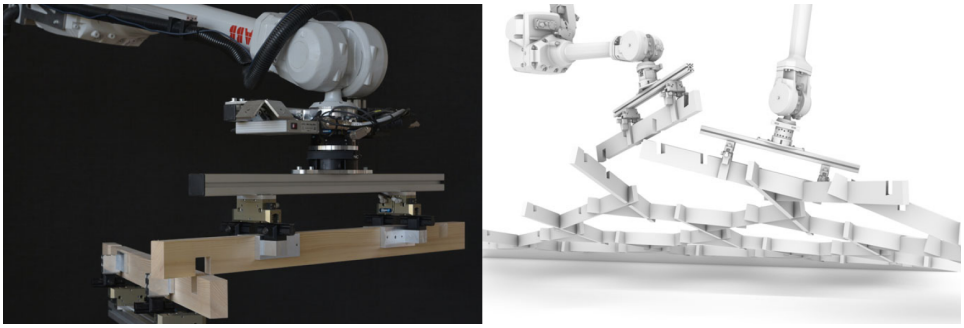

Figure 2.4: Tightening timber joints with robotic arms [Apolinarska et al., 2021].

research by a team from ETH has developed an autonomous robotic excavation system for building dry stone structures on site [Johns et al., 2023] (left in Figure 2.5). This research modifies an off-the-shelf robotic excavator with a multi-sensor suite of LiDARs and cameras, enabling the robot to autonomously pick and place stones in a predefined pattern on an actual construction site.

In the industry, BUILT Robotics in San Francisco leverages advances in AI/ML to enhance existing large construction vehicles by developing AI/ML powered autonomous upgrade kits [Built Robotics, 2024]. Boston Dynamics has developed a series of mobile robots for construction related tasks, including the Spot robot for site inspection [Halder et al., 2023] and the Stretch robot for material handling [Ackerman, 2022]. Dusty Robotics has created a robot for layout and surveying tasks on construction sites, aiming to improve accuracy and efficiency in these critical construction processes [Lau, 2020] (right in Figure 2.5). These initiatives underscore the growing interest and investment in robot-centric research, reflecting the potential of robots in transforming certain sectors of construction practices and workflows.

While the predominant robot-centric trend has seen a growing body of interdisciplinary research and practical on-site applications, several unaddressed concerns persist, particularly regarding the integration of these

Figure 2.5: Field mobility robots for autonomous excavation [Johns et al., 2023] and layout and surveying [Lau, 2020].

technologies into existing construction workflows and human work practices. Firstly, the current state of AI and robotics has not achieved the level of expected autonomy and dexterity. Despite its promise, significant efforts are required before highly adaptive, safe, and deployable autonomous robotic dexterity in complex work environments like construction can be realized. Secondly, construction sites, whether on-site or off-site, are inherently more complex than standardized factories or organized warehouses. They are bespoke and dynamic, constantly requiring manual, situational and contingent improvisation. The repetitiveness of construction tasks does not equate to the predictability and reproducibility of typical manufacturing pipelines. Skillful, dexterous, and highly adaptive human work is an essential, always-already presence inherent to real-world construction. Overestimating robotic autonomy or oversimplifying human factors and roles in existing realities risks misunderstanding the unique challenges of construction environments. In order to overcome these barriers, it is essential that we candidly recognize the technical limitations, such as adaptability, error handling, and customization, of solely relying on robotic systems to resolve all construction challenges, and instead, factor in the underappreciated existing human skills and practices.

2.1.3 Design-Centric Trend

Contrasting with the material-centric and robot-centric trends, a distinct group of researchers advocates a design-centric approach. Emphasizing the intricate and customized nature of construction, this perspective prioritizes simplifying construction at its inception—the design—by developing a new sets of building syntax [Retsin, 2019b]. The primary motivation is to frame the issues in automating construction as a design challenge rather than solely a technological hurdle, aggregating and merging complexities associated with existing construction tasks through a simplified, redesigned system. This approach also critically leverages robots, but more as *executors of the syntax*. Interdisciplinary methodologies from discrete geometry, shape grammar, physical grammar, and design are employed to rethink the fundamental systems of physical components in construction.

Representative of this trend, researchers from UCL Bartlett's Design Computation Lab propose modular, predefined building components and their associated repetitive tectonic systems [Retsin, 2019a] (Figure 2.6). This approach shifts customization from the traditional focus on blocks, materials, procedures, or robots to the unique assembly of a set of standardized components. Some representative projects by the group, such as the Unit 19 and Semblr projects, demonstrate the use of discrete automation frameworks to reimagine architectural scales and scenarios. The Semblr project, led by Ivo Tedbury[3], reconfigures Le Corbusier's Maison Dom-Ino using discrete parts, uniting design, material, robotics, and construction under a novel syntax of spatial configuration. These discrete elements serve both as material for participatory design and as modular units for robotic end-effectors. In the Chamfer project, connections between fragments are organized in a matrix, showcasing the system's capacity to accommodate

[3]A detailed visual account of the project is accessible here: `https://www.tedbury.co.uk/semblr`

complexity and variation [Claypool, 2019]. While these conceptual prototypes are primarily demonstrated at smaller scales in labs or through manual student construction, their potential remains significant [Retsin, 2019b].

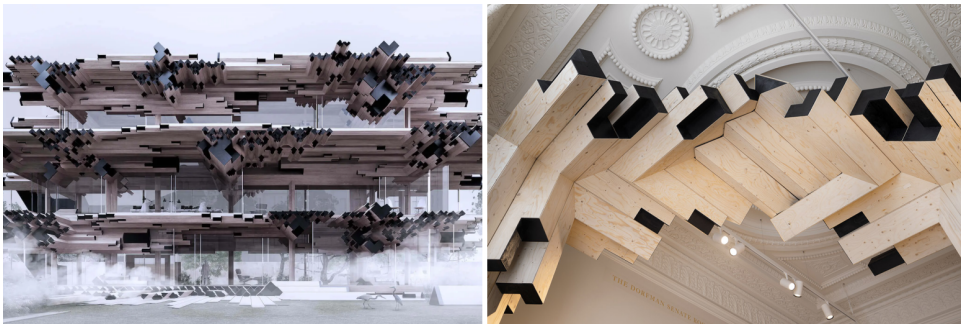

Figure 2.6: Design-centric research in construction robotics, featuring representative installations by Gilles Retsin from UCL Bartlett [Retsin, 2019b]. Image credit: *Gilles Retsin interview, DesignWanted*, source: `https://designwanted.com/gilles-retsin-interview/` (left), and *Real Virtuality, Gilles Retsin Architecture*, source: `https://parametric-architecture.com/real-virtuality-gilles-retsin-architecture/` (right).

Despite its innovative and reformative outlook, the design-centric trend still faces challenges in real-world application, particularly concerning how the proposed building syntax can be introduced and integrated into existing and established construction conventions and practices. It remains questionable whether the limited set of lab-curated, predefined building parts can be broadly accepted in diverse real-world scenarios and whether it has the capability to handle the various bespoke needs present in the real world. Additionally, it is uncertain whether such a simplified building syntax might restrict spatial creativity, making certain building options infeasible due to the limitations of the proposed syntax. Lastly, while the assumption that design simplification may lead to fewer construction links, the practical challenges of robotically realizing the intended syntax at a real-world building scale present unaddressed complexities and challenges.

2.1.4 The Emerging Human-Centric Trend

Aside from the previously mentioned trends, a human-centric trend in construction robotics is emerging. This trend acknowledges the unique capabilities of both humans and robots, aiming to leverage their respective strengths to enhance construction processes. Unlike the robot-centric trend, which focuses on replacing human labor with autonomous robots, the human-centric trend aims to directly augment human capabilities through recent robotic technologies.

For instance, one aspect of this trend involves developing exoskeleton systems to enhance human's physical capabilities and reduce the risk of injury [Cho et al., 2018]. Researchers at Harvard's Wyss Institute have developed a soft exosuit that reduces the metabolic cost of walking and running, enabling workers to carry heavy loads with less effort [Nuckols et al., 2021] (left in Figure 2.7). This exosuit is designed to be lightweight, flexible, and comfortable, allowing workers to move freely while providing necessary support. Matsubara et. al at Kyoto University and Nara Institute of Science and Technology contributed to this theme of research by leveraging direct physical user interface to achieve more comfortable and safe soft acuator models and movement assistance [Hamaya et al., 2021]. In these examples, exoskeletons commonly are designed through a combination of soft materials and actuators to provide support during physically demanding tasks.

Besides augmenting workers' physical capabilities, some research explores how robots can assist workers in performing complex, skillful tasks through various interfaces. For example, researchers at Stanford University have developed a haptic/UI robotic system that assists workers in manual tasks such as welding and joint sealing [Brosque et al., 2021] (right in Figure 2.7), reducing physical strain and cognitive load. The system uses a combination of computer vision, machine learning, and robotic manipula-

tion to guide workers in conducting skillful tasks accurately and efficiently.

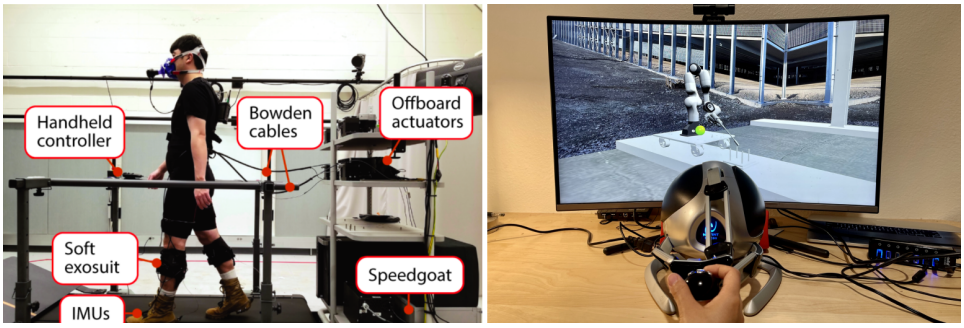

Figure 2.7: Examples of human-centric research in enhancing physical work capabilities, including exoskeletons [Nuckols et al., 2021] (left) and haptic/UI robotic systems [Brosque et al., 2021] (right).

This research sits in and aligns with the human-centric trend. However, rather than depending on explicit interfaces or reconfiguring existing construction workflows, it focuses on developing social robots that integrate into current workflows, providing semi-autonomous or autonomous support. The objective is to create robotic systems that assist human workers, instead of fundamentally changing their established practices. This approach recognizes the distinct capabilities of both humans and robots, aiming to harness their combined strengths to improve construction processes.

2.2 Developing Socially-Aware Navigation Robots

In the realm of robotics, there is a growing focus on developing mobile robots capable of operating in human-populated environments. This field, commonly referred to as social robot navigation, has evolved considerably over the past decades [Mavrogiannis et al., 2021]. Initially exemplified by museum-guiding robots like RHINO [Burgard et al., 1999] and MINERVA [Thrun et al., 1999], the field has transitioned from viewing people merely as dynamic obstacles to considering them as active, behaviorally responsive

agents. Robots' abilities have also progressed from basic collision avoidance to encompassing a broader understanding of social norms and behaviors. Current expectations for robots include avoiding disruption of social groups [Katyal et al., 2022], approaching individuals in a non-intrusive and courteous manner [Kato et al., 2015], and respecting personal space in a human-like fashion [Torta et al., 2013, Montesdeoca et al., 2017]. The practical application of these developments spans various sectors, including Ambient Assistive Living (AAL) for the elderly [Calderita et al., 2020], public space delivery and services [Kannan et al., 2021], and interactive therapy and education for children [Calderita et al., 2014, Escobar-Planas et al., 2022].

Recent developments, especially in Reinforcement Learning (RL), have significantly expanded the research possibilities in this field. Notable contributions from MIT researchers Chen and Everrett, like SA-CADRL [Chen et al., 2017] and GA3C-CADRL [Everett et al., 2018], have opened up promising research directions. By framing the decision of the robot's next action as a Markov Decision Process (MDP) with observable states and a reward function, deep reinforcement learning has proven effective in enabling robots to predict and respond adaptively to dynamic human activities, even in dense crowds. Subsequent studies have expanded this research by employing various recent deep neural network architectures and addressing specific aspects of social awareness and compliance. For instance, researchers from UIUC leveraged a decentralized structural-Recurrent Neural Network (DS-RNN) architecture to investigate the spatial and temporal dependencies of human trajectories in crowded environments [Liu et al., 2021]. Another group of researchers from Stanford University and Berkeley proposed a novel pipeline accentuating the dynamic relational reasoning among multiple agents [Li et al., 2024].

Despite the promising results these studies have shown in tackling in-

creasingly dense human crowds and complex social interactions, several inherent limitations arise when applying these methods to cluttered, populated work environments. Firstly, these studies often assume a relatively open, uncluttered environment, such as open squares, wide corridors, or empty classrooms, where the robot's primary goal is to navigate smoothly and efficiently through the crowd. This assumption is closely related to the initial research intent of developing robots to navigate in densely crowded public spaces. However, to my knowledge, the performance of these methods in cluttered indoor construction site settings has not yet been examined. Secondly, due to the inherent assumption of viewing people mostly as goal-reaching individuals or grouped pedestrians, the behavioral norms of these agents in testing these algorithms are largely monolithic. When applied to a work environment, there is an additional layer of contextual work dynamics and social norms displayed by workers. The contextual layer is highly associated with work activities, roles, and contents. This domain difference is less explored in the current research landscape yet is a critical, real-world challenge when developing socially aware companion robots for active workers. Addressing such distinctions and challenges requires not just a single algorithmic effort but an entire framework and pipeline, including systematic customization and redesign of experiment setups, data collection, algorithmic alignment and fine-tuning, and context-specific evaluation metrics.

With these prospects and limitations of the key algorithmic advances in social robot navigation as backbone, the following sections cover some supplementary technical aspects in practically and physically developing socially aware navigation robots. Unlike testing an algorithm in a symbolic simulator, where the input of human and robot states is easily supplied by numeric calculations, developing physical robots that can showcase social awareness among people in the real world also requires capabilities like site perception, people detection and tracking, and hierarchical path planning.

While these capabilities are less mentioned and often assumed to be a basic prior in social navigation research, they involve significant engineering efforts and present unique challenges in particularly cluttered, dynamic, and human-populated work environments.

2.2.1 Site Perception

For the robot to reach a goal position, it must first perceive its environment. This involves two key processes: mapping and localization. The mapping process involves reconstructing a representation of the environment, whereas localization determines the robot's real-time position within the map. In the context of social navigation, the robot must perceive not only the physical layout of the environment but also the presence and movement of people within it. In simpler environments, such as spacious and structurally static areas, this can often be addressed using 2D scan-based methods and toolkits like gmapping and AMCL. However, as the environment becomes geometrically complex and non-stationary, the challenges of both mapping and localization increase. Consequently, methods leveraging various sensors, such as cameras, LiDAR, and depth sensors, emerge to provide the robot with a comprehensive understanding of its surroundings.

Recent advancements in SLAM (Simultaneous Localization and Mapping) algorithms have enabled robots to create detailed maps of their environment in real-time, allowing them to navigate autonomously and adapt to changes. For instance, Lidar Odometry and Mapping (LOAM) [Zhang and Singh, 2014] and its successive methods [Shan and Englot, 2018, Shan et al., 2020] leverage the dynamic matching of point clouds to create detailed 3D maps of the environment, while enabling robots to localize themselves by providing high-quality odometry estimates in complex, cluttered environments. Visual SLAM methods, such as ORB-SLAM [Mur-Artal et al.,

2015] and its successors [Mur-Artal and Tardós, 2017, Campos et al., 2021], offer an alternative approach to mapping and localization by using visual features extracted from camera images. These methods have been successfully applied in various robotic applications, including autonomous driving, warehouse automation, and service robots. However, the challenges of mapping and localization in dynamic, cluttered environments like construction sites remain a significant area of research, requiring further investigation to develop robust and reliable solutions.

2.2.2 Detection and Tracking of People

Another critical capability of socially aware robots is the detection and tracking of people within their operating environment. Depending on the characteristics of the context and the choices of hardware, various sensors can be used for this purpose, including cameras, LiDAR, depth sensors, and thermal sensors. Primarily, LiDAR-based and camera-based methods leveraging various computer vision techniques and deep learning algorithms have shown promising results in detecting and tracking people in real-time. For instance, YOLO (You Only Look Once) [Redmon et al., 2016, Jiang et al., 2022, Wang et al., 2023] is a popular vision-based deep learning model used for real-time object detection, including people, in images and videos. YOLO's single-stage detection approach enables fast and efficient detection of people, making it suitable for real-time applications. PointPillars [Lang et al., 2019] and PointNet [Qi et al., 2017] are examples of LiDAR-based deep learning models used for 3D people detection, in point cloud data. There are also hybrid methods that combine multimodal information from multiple sensors to improve detection accuracy and robustness. For instance, researchers at MIT have developed a multimodal fusion model, BEVFusion, that combines camera and LiDAR data to synthetically detect and track

objects in bird-eye-view [Liu et al., 2023].

In addition to detection, tracking people across multiple time frames is essential for maintaining a consistent understanding of their movements and interactions. Tracking algorithms, such as Kalman filters [Welch et al., 1995] and Particle filters [Arulampalam et al., 2002], have been widely used for this purpose. However, recent advancements in deep learning have led to the development of deep learning-based tracking algorithms that offer improved accuracy and robustness. For instance, Deep SORT (Simple Online and Realtime Tracking) [Wojke et al., 2017] is a deep learning-based tracking algorithm that combines object detection and tracking to maintain consistent tracks of people across multiple frames. Similarly, FairMOT (Fair Multiple Object Tracking) [Zhang et al., 2021] is a deep learning-based tracking algorithm that leverages a combination of object detection and re-identification to track people in crowded scenes. These algorithms have shown promising results in tracking people in complex, cluttered environments, particularly with a certain degree of occlusions.

2.2.3 Path Planning

While social robot navigation primarily focuses on local path planning in human-populated environments, the global path planning stage, which determines the robot's optimal path from its current position to the goal position, especially over long distances, is equally crucial for ensuring overall safe and efficient navigation. In the context of construction sites, global path planning presents unique challenges due to the dynamic nature of the environment, the presence of obstacles, and the need to avoid disrupting ongoing construction activities. Traditional path planning algorithms, such as A* [Hart et al., 1968] and Dijkstra [Dijkstra, 2022], although efficient, often struggle in dynamic environments or when human workers are present.

Recent advancements in global path planning algorithms, such as RRT* [Karaman and Frazzoli, 2011] and RRT-Connect [Kuffner and LaValle, 2000], have shown promise in addressing these challenges by leveraging probabilistic sampling and tree-based search techniques to efficiently explore the robot's configuration space and generate collision-free paths to the goal position. Beyond the efficiency of these algorithms, integrating the global path planning algorithm with the local path planning algorithm poses additional challenges. Key concerns in this integration include the mechanism of discretizing and sampling the action space, how such discretization may affect subsequent probabilistic methods, and the real-time computational acceleration of the overall path planning pipeline. This requires a deep understanding of the assumptions and priors of these algorithms, going beyond simply comparing standalone navigation performance. It is essential to ensure the appropriate integration of the two stages of path planning.

2.3 Reinforcement Learning for Adaptive Robots

The role of artificial intelligence (AI), specifically Reinforcement Learning (RL), is crucial in the evolving landscape of robotics and machine learning. RL, a subset of AI and Machine Learning (ML), focuses on how agent learns to act in an environment by maximizing a cumulative reward [Sutton and Barto, 2018]. This learning paradigm, distinct from traditional methods like expert-supervised solutions or hand-crafted models, involves the agent developing an adaptive decision-making policy through trial-and-error interactions within their environment [Sutton and Barto, 2018]. The advent of Deep Reinforcement Learning (DRL), exemplified by its success in AlphaGo [Silver et al., 2017] and Atari player [Mnih et al., 2013], has sparked growing interest in applying RL to a variety of robotic control tasks. These tasks include contact-rich manipulation [Wu et al., 2021, 2022, Elguea-Aguinaco

et al., 2023], humanoid bipedal walking [Rodriguez and Behnke, 2021, Li et al., 2021], autonomous car driving [Sallab et al., 2017, Kiran et al., 2021], and drone flight control [Song et al., 2021, Fang et al., 2020, Kaufmann et al., 2023]. RL's suitability for these tasks stems from the dynamic and complex nature of real-world robotic behaviors, which often defy explicit modeling or detailed manual scripting.

2.3.1 Reinforcement Learning in Construction Robotics

Recently, RL's applicability has extended to the construction sector [Delgado and Oyedele, 2022]. For example, in the ReAC group, we proposed a distributed multi-drone framework using multi-agent reinforcement learning methods for semi-autonomously performing collective tasks such as bricklaying and spray coating [Fang et al., 2020]. Researchers from ETH utilized a distributed RL method to automate the timber joint assembly task [Apolinarska et al., 2021]. Based on state input extracted from real-time force/torque and camera data, the method trains a policy to enable adaptive insertion between two etched beams. However, a common characteristic in these examples is the primary interaction of robots with materials or other robots, largely overlooking the coexistence of human workers in the same environment. For instance, in our multi-drone research, the group specifically developed a system shutdown protocol for occasions of human interventions and supplies, where workers may need to enter the drone's operation space. Consequently, RL's adaptability has been mainly harnessed to improve robot-material and robot-robot proficiency and dexterity, leaving its potential for on-site mobility among workers, an actual and realistic presence, a largely untapped area of investigation. This observation highlights a significant yet underexplored area: the application of RL for adaptive robot mobility in real-world dynamic construction environments, specifically those

actively occupied by busy workers.

2.3.2 Different Types of RL Algorithms

To meaningfully delve into the technical aspects of RL, an overview of various RL algorithms is provided here. While there are multiple ways and dimensions to categorize RL algorithms, this review primarily focuses on the categorization relevant to this research: value-based, policy-based, and actor-critic methods.

Different categories of RL algorithms have their unique strengths and weaknesses in practice, making them suitable for different types of problems and environments. Value-based methods, such as Q-learning [Watkins and Dayan, 1992] and Deep Q-Networks (DQN) [Mnih et al., 2013], estimate the value of taking a particular action in a given state. Policy-based methods, such as REINFORCE [Williams, 1992] and Proximal Policy Optimization (PPO) [Schulman et al., 2017], directly learn a policy, a probabilistic function, that maps states to actions. Actor-critic methods, such as DDPG [Lillicrap et al., 2015] and TD3 [Fujimoto et al., 2018], combine the advantages of both value-based and policy-based methods by learning both a policy and a value function.

The previously mentioned group of social robot navigation algorithms commonly leverages value-based or actor-critic RL algorithms. This choice is related to the nature of the problem and the use of discrete action spaces. Rather than having the RL algorithm directly output an action from a continuous action space, the RL algorithm ranks and selects from a finite set of actions. This seemingly minor distinction, in practice, can affect the robustness of the robot's performance. This setting, in turn, is particularly suitable for employing value-based methods, where the robot can choose the

action with the highest value. Compared to policy-based methods, value-based methods are also less sensitive to hyperparameters, which further contributes to stability. In this research, the mechanism of ranking among discrete action spaces offers an additional advantage when working with other layers of robot path planning, where the discretization and sampling of the action space is a common practice.

2.3.3 Simulation Environment for RL

A critical link in training and developing RL algorithms for robotic applications is the simulation environment. Simulators provide a controlled, safe, and scalable virtual platform for training and testing RL algorithms before deploying them in the real world. In the domain of social robot navigation, efforts have been made to develop various training and evaluation utilities. As a starting point, researchers from MIT developed the `gym-collision-avoidance` environment (GCA) [Chen et al., 2017], which captures dynamic agents as symbolic circles and offers a preliminary set of policies to simulate different kinds of collision avoidance behaviors. Extending this effort, researchers from CMU further enriched the simulation environment with more diverse and customizable options, enabling more flexibility to experiment with various agent behavioral policies [Shum et al., 2022]. While both environments primarily leverage the `TensorFlow` [Abadi et al., 2015] libraries, later efforts of `CrowdNav` [Chen et al., 2019] redeveloped a similar software utility using `PyTorch` [Paszke et al., 2019], making it more succinct, modular, and accessible. Seemingly simplistic and visually abstract, these simulation environments provide a solid foundation for training and evaluating algorithms with a focus on diverse agent behaviors and policies.

In addition to these symbolic and abstract simulation environments,

there are also more photorealistic 3D environments. NVIDIA Omniverse Isaac [NVIDIA, 2023] and Unity ML-Agents [Juliani et al., 2018] are general-purpose 3D simulation platforms that include a variety of assets for intuitive GUI customization. However, creating a simulation environment leveraging these platforms that can capture the desired characteristics of this research is an ambitious undertaking. While these platforms offer rich libraries of static assets, animating them in a lifelike, contextually detailed fashion remains a challenge. Additionally, there are certain software bottlenecks in the current versions of these platforms, and the room for deep customization underneath the hood is limited.

For more dedicated 3D environments, researchers from Yale University have open-sourced two versions of SEAN (Social Environment for Autonomous Navigation) [Tsoi et al., 2020, 2022], a photorealistic environment that combines 3D scene objects, human avatar meshes, and available sets of algorithms deployable to the avatars. SEAN offers a useful benchmarking platform for evaluating algorithm performance in more realistic and visually appealing scenarios, particularly for recent endeavors in end-to-end technical stacks and methods. However, the available scenarios do not currently include construction sites, and the customization of more sophisticated avatar behaviors is a significant technical barrier. For representing the context-specific nuances of worker activities, described in more detail in later chapters, 2D symbolic environments are surprisingly more adaptable and effective than currently available 3D environments.

As a summary of this chapter, several key insights emerge. Firstly, the current landscape of construction robotics encompasses multiple trends, such as material-centric, robot-centric, and design-centric approaches. The emerging human-centric trend, in particular, underscores the importance of considering human factors in the development of robotic systems for real-

world construction work. Secondly, with advances in social robot navigation, there is a growing body of research in key algorithms and supplementary technical aspects, such as site perception, people detection and tracking, and hierarchical path planning. These advancements provide a promising foundation for developing socially aware work companion robots in existing human workplaces, particularly in construction contexts, where this approach remains unexplored. Thirdly, ongoing research in Reinforcement Learning (RL) for adaptive robots, although previously focused on construction tasks involving materials or other robots, presents a unique advantage for enabling adaptive robot navigation in worker-robot interactions and on-site navigation among workers. These insights collectively profile and locate a promising yet challenging landscape for this research, which will be detailed in the following chapters.

2.4 Context of Carpentry Formwork

In this section, I elucidate the process of carpentry formwork observed during the rough carpentry stage on-site, establishing the grounding context for this research. The descriptions and insights draw and extend from a concurrent qualitative site study conducted by the ReAC research group[4]. This study spanned a 15-week observation period at a multi-level building, where research group members employed ethnographic methods to qualitatively document and analyze eleven construction tasks identified within what is often referred to as the "rough carpentry stage". Leveraging critical insights from these tasks, my colleagues and I explored potential robotic support capabilities and functions, considering both qualitative and technical aspects. Our focus was on (1) identifying fundamental functionality

[4]A detailed discussion of the qualitative aspects of this study can be found in Emek Erdolu's PhD dissertation Erdolu [2024].

challenges, (2) enhancing human-robot interaction, and (3) aligning with our group's technical capabilities. Throughout our discussions, certain activities and support opportunities were deemed unsuitable. For instance, we decided against pursuing floor cleaning in fireproofing, because extensive efforts will inevitably be devoted to developing a debris removal and suction robotic hardware system, away from the focus of human-robot interaction itself. We also discarded the idea of robotic support for spray coating due to the inherent requirement for the robot to operate at varying heights with continuous paint supplies, a capability beyond the reach of our current robotic resources.

The group discussion and analysis ultimately pinpoint carpentry formwork (Figure 2.8) as a work activity potentially well-suited for robotic support, based on three primary insights. First, the workflow is inherently collaborative, with two workers frequently cooperating and exchanging tasks. This provides an opportunity for studying the dynamics of labor-intensive human work, particularly how robots can integrate into such interactions. Second, the demands of tool/material transportation are notably high, although often overlooked by workers who assume this physical strain is integral to the job. This qualitative observation of frequent heavy load bearing across the site is also substantiated by statistics gathered from site observation. Third, compared to precision-oriented tasks like welding or painting, the facilitative tool/material supply task in this workflow is more forgiving in terms of precision requirements. Instead, it focuses on other aspects of human-robot interaction, such as social awareness and comfort around workers. This allows for a more relaxed and adaptable robotic support system, enabling the research to focus on the aspiring human work fabrics itself, rather than strictly on improving a specific robotic task's precision. These insights, combined with the potential for a significant impact on workplace safety and workflow fluency, lead to the selection of carpentry formwork as

the specific context of this research.

In more depth, carpentry formwork is an essential process that precedes concrete pouring in on-site construction. It mostly involves a two-person team working together to install panels around tied rebar cages, creating the external molds required for concrete structures. The activities observed by our research group predominantly involved setting up aluminum panels to form a future concrete wall. For consistency within discussions and documentation, I refer to the primary worker as Worker A and the assistant as Worker B, aligning with the terminology used by our group. Worker A's primary responsibility is installing mold panels at specific locations using various tools and hardware. Worker B plays a supportive role, handling all preparatory tasks such as arranging the necessary panels, hardware, and tools for Worker A (see Figure 3.1, Figure 3.5 and Figure 3.6). A significant portion of Worker B's tasks involves the frequent transportation of tools and hardware to the current work zone from either a central workbench or another work zone. As observed by my colleague Erdolu [2024], there are instances when Worker B may need to make up to around 19 trips of 15∼20m within an hour to support the workflow. This intensive and frequent movement, as well as the physical demand of carrying heavy loads of tools and hardware, highlight prime opportunities for integrating a navigation robot to serve as an additional assistant in the process.

Parallel to the selection of carpentry formwork as the research context and based on preliminary findings and discussions, the group identified unmanned ground vehicles (UGVs) as the primary type of robot, as opposed to drones or fixed robotic arms. While the robot base and model are actively iterated upon (see Chapter 4 for detailed iterations of the robot base and chassis) and robot support scenarios are being developed, the candidate scenarios and functions inevitably require some degree of on-site mobility,

Figure 2.8: Carpentry formwork activities observed on-site. Image credit: ReAC Research Group, 2022, diagram and photo by Emek Erdolu Erdolu [2024].

load-bearing, and navigation-related features operable on a rough flat surface. These features fit the profile of a mid-sized, rover-like UGV [5]. Consequently, this choice also influenced the development of the envisioned support scenarios, particularly concerning the technical capabilities this type of robot may be customized to perform.

In the upcoming discussion, I will explore the potential robot support scenarios and functions that have emerged from the observed opportunities. With the envisioned robotic support in mind, the context of carpentry formwork will be scrutinized from two perspectives: (1) site and material, and (2) work activities. Building upon the qualitative research conducted by the ReAC research group, my analysis extends this effort by approaching it through a technical lens with a focus on robot navigation. The distinct specificity of carpentry formwork is insightfully brought to the forefront and emphasized via methods of decomposition, abstraction, and categorization. This approach not only facilitates a deeper understanding of material-specific and behavior-specific challenges but also lays the groundwork for identifying high-level, complex technical difficulties introduced by these context-specific nuances. Overall, this analysis clarifies the critical rationale behind the technical decisions documented in the ensuing chapters.

[5]This finding was reached after testing different platfoms, and was also influenced by their availability to the research group.

Chapter 3

Hypothesis

This research proposes that Reinforcement Learning (RL)-based social navigation methods, when adapted and integrated into a comprehensive robotic framework, can effectively support on-site construction workers by enabling safe, adaptable, and socially aware robot delivery and companion functions. However, rather than simply applying and merging existing social navigation algorithms, this research must also address the unique social and material aspects of workers' activities and work environments, which differ significantly from the assumptions made in previous social navigation studies. A key technical challenge lies in developing an algorithmic methodology that not only leverages advancements in existing research but also systematically addresses the navigation challenges specific to this context. In more detail, the hypothesis is grounded in tackling the following sub-technical challenges:

- **Mid-Distance Robot Navigation on Complex and Populated Construction Sites:** The hypothesis posits that RL-driven hierarchical motion planning, when combined with a robust robot design and hardware setup, can achieve safe and efficient navigation on construction sites characterized by dynamic environments and obstacles. The

system must accommodate the unpredictability and clutter inherent to these sites while maintaining stable navigation paths.

- **Accurate Profiling and Mapping of Unstructured Site Conditions:** facilitating adaptive robot behavior, the framework must employ advanced unstructured site perception and robot state estimation techniques. These methods are hypothesized to enable real-time, accurate mapping and profiling of the construction environment, which is essential for robot localization and navigation.

- **Real-Time Detection and Interpretation of Worker Activities:** This research further hypothesizes that an integrated worker detection and tracking pipeline, utilizing a customized combination of sensors and computer vision methods, will enable the robot to detect and interpret worker activities and movements in real-time. This capability is crucial for accurately providing and estimating the state information of nearby human workers' trajectories. Such ability to distinguish and track workers in a complex work environment forms a critical foundation for effective robot assistance, enabling prompt navigation responses when human workers enter the robot's operational vicinity.

- **Adaptive and Socially Aware Robot Navigation and Interaction:** By integrating the aforementioned software components into a cohesive RL-driven robotic framework, the system will also undergo contextual algorithm fine-tuning to account for the uncertainty and variability of worker activities highlighted in this research. The goal is to provide comfortable, tangible, and adaptive assistance tailored to the specific needs and realities of the observed group of workers. The hypothesis posits that the robot will be capable of navigating safely and comfortably around workers in close proximity, offering support

through functions such as sending, summoning, and accompaniment (discussed in more detail in later sections).

- **Physical Robot Design and Prototyping:** The hypothesis importantly includes the development of a physical robot prototype capable of carrying tools and hardware, incorporating various sensors and onboard computing, being safe to operate around workers, and maintaining a stable structure for robust navigation on rough, flat surfaces. The robot design will be informed by the specific site and material conditions, as well as the above algorithmic and technical requirements. The robot prototype acts as a physical realization of the framework and serves as a testbed for evaluating the proposed robotic support functions, both in a lab setting and on real construction sites.

3.1 The Robot Support Scenario

Through iterative discussion and brainstorming, our group identifies two principal robotic support scenarios aimed at mitigating the burdens associated with site traversal and the physical demands of transporting materials. These discussions, which mostly take place in the initial stages of technical development, ultimately lead to the definition of three essential robotic support functions: **send**, **summon**, and **accompany**. The primary objective behind the proposed scenarios and functions is to reduce the physical strain on Worker B, thereby allowing him to focus on his primary task of panel preparation—a crucial and detailed task critical to the carpentry formwork process. These functions, detailed below and illustrated in Figure 3.1, include:

- **Sending** (see Figure 3.1): This function is utilized when tools or hardware are needed from a distant location, such as a central work-

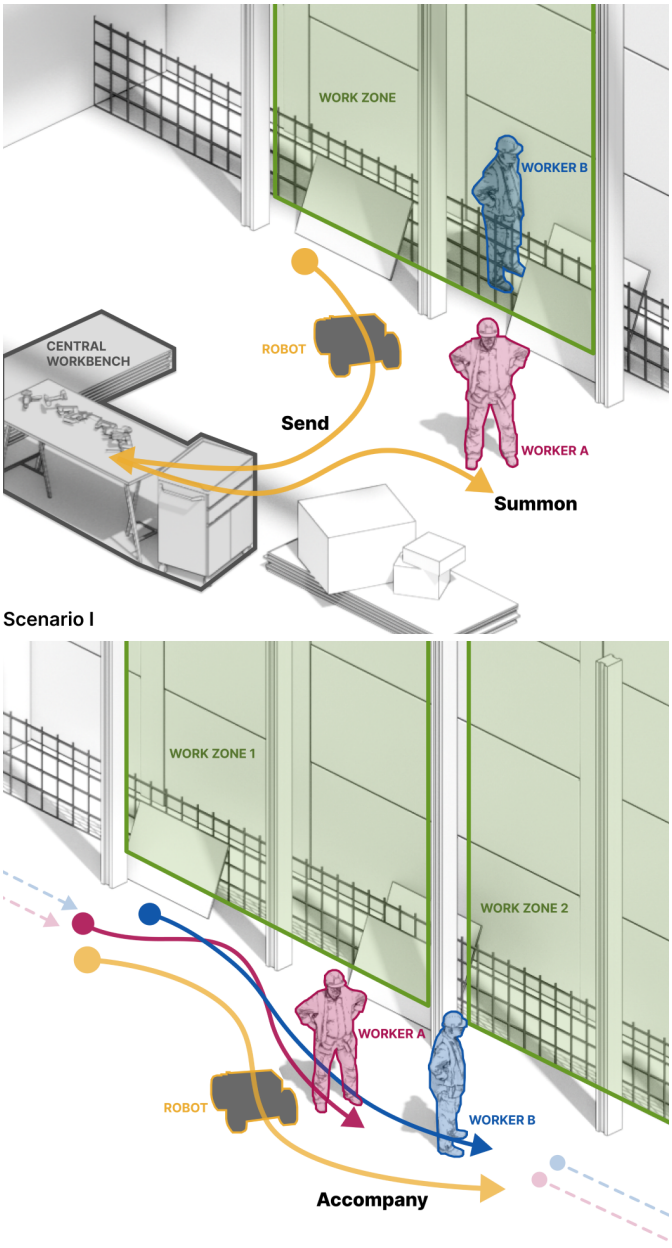

Figure 3.1: Robot support scenarios and functions: sending, summoning, and accompanying [Wu et al., 2024b].

bench. Rather than halting ongoing tasks, Worker B can command the robot to autonomously fetch the required tools or materials using a handheld device. The robot navigates to the workbench, where a designated worker managing site tools prepares and loads the robot with the necessary items. The robot then delivers the items directly to Worker B, allowing for seamless continuation of the panel preparation process without disrupting workflow.

- **Summoning** (see Figure 3.1): When specific tools or materials are needed at the work zone, Worker B can summon the robot back to their location after it has been sent to retrieve items. With a simple command, the robot promptly returns to the work zone, delivering the required items. This function enhances the efficiency of the workflow by minimizing downtime and ensuring timely delivery without requiring Worker B to leave the task at hand.

- **Accompanying** (see Figure 3.1): In this function, the robot serves as a mobile support unit, akin to an intelligent tool cart that accompanies workers as they move between work zones. This feature is particularly valuable during transitions between tasks or locations, as it reduces the need for workers to manually carry heavy tools or materials. The robot follows and stays in close proximity to workers, providing continuous support as they progress through different stages of the work process.

The proposed functions and scenarios demonstrate a strategic approach to integrating robotic support within the existing carpentry formwork process, with a focus on enhancing worker safety and comfort while reducing repetitive physical strain. Rather than displacing human roles, these auxiliary supports are designed to complement the workforce by taking on physically demanding, peripheral tasks. This approach allows carpentry workers to concentrate on the more skilled aspects of their work, fostering a work

environment where the robot's capabilities are aimed at supporting, not replacing, human efforts.

3.2 Specificity of site and material

While the group's robot base availability and capability influenced the selection of the site—primarily flat but cluttered rough indoor surfaces—the realization of the envisioned robotic support necessitates a more thorough understanding of the physical environment in which the robot will operate. Identifying the specific physical conditions relevant to robot navigation is essential and will provide critical guidance for future technical decisions concerning robot sensing and planning. In addition to understanding the site conditions and objects, it is imperative to closely examine the tools and hardware that carpentry workers commonly transport between the workbench and their work zones. Estimations of the weight and dimensions of these items will considerably impact decisions related to the robot's tool container and its space and load capacity.

3.2.1 Site Objects and Conditions

Construction sites during the hand-detailing and interior installation phase are notably messy and cluttered, in stark contrast to the well-arranged and organized settings found in other manufacturing environments. Interior walls and structures stand bare and incomplete. Materials, tools, and machinery related to the work are dispersed and piled within work zones for convenience and temporary access. Dust, debris, and residues from ongoing activities accumulate on the floor, awaiting subsequent cleanup. Amidst this complexity, I aim to categorize the various objects by examining and

analyzing time-lapse camera footage[1]. The focus of the investigation is on areas pertinent to carpentry formwork activities. This chiefly encompasses the linearly arranged work zones adjacent to the target formwork wall, the workbench area, and the intermediary spaces through which the robot is expected to navigate frequently. I group the objects in this area into four categories, with an emphasis on their 3D geometric features. Recognizing the unique attributes of each category will aid in the development of tailored technical solutions and offer a practical guideline for later conducting physical experiments in a laboratory sandbox environment.

- **Category I - Low, Temporary, Shape-Varying Objects**: This category encompasses items such as (1) temporary material piles, (2) hand tools left on the ground, (3) material wastes and large residues from ongoing work, (4) wires, cables, and extension sockets for electric tools, and (5) poles, warning signs, and temporarily positioned long timber for delineating areas. These objects, frequently scattered and irregularly positioned, present a navigation challenge due to their highly transient nature and the low-lying, small-scale, fragmented geometric shapes they possess.

- **Category II - Low to Medium Height, Bulky Objects**: This category includes objects such as (1) large trash bins, (2) stacked boxes, (3) material piles stacked on wooden pallets, and (4) containers. Throughout the duration of a working day, their positions remain relatively unchanged, adjusting only occasionally according to the requirements of the work. Geometrically, these objects typically present as large, voluminous entities that the robot should navigate around.

- **Category III - High, Fixed Installations**: These permanent fix-

[1]The studied time-lapse recording comes from the ReAC Research Group's on-site observation and data collection.

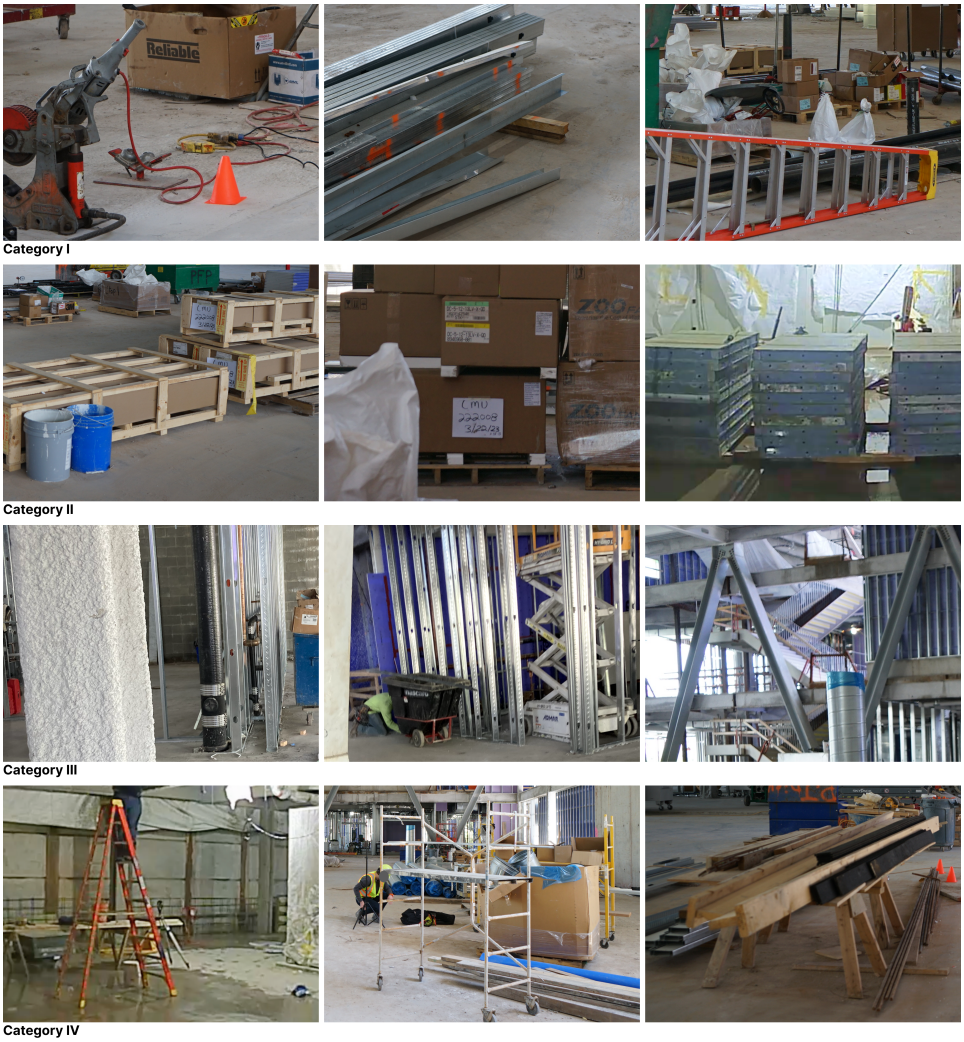

Figure 3.2: Four key categories of site objects, (1) low, temporary, shape-varying objects, (2) low to medium height, bulky objects, (3) high, fixed installations, and (4) high, thin, temporary objects.

tures on construction sites encompass structures like (1) steel columns, whether decorated or plain, (2) walls, both fully erected or partially built, (3) concrete structures, (4) rebar cages, (5) concrete mold panel structures, and (6) vertical pipes and ducts. The unfinished nature of these structures—characterized by the slender and curved profiles of exposed steel columns, the dense clusters of thin rebars and steel

beams, and the irregular configurations of partially completed walls—distinguishes them from their fully constructed counterparts. This distinction poses unique and nuanced navigational challenges for robot sensing, necessitating a cautious approach due to their distinct geometries and states of completion.

- **Category IV - High, Thin, Temporary Objects**: This category comprises items such as (1) lifter carts, (2) twin-step ladders, both open and closed, (3) four-legged tables and working surfaces, (4) high steel storage racks and (5) heavy-duty tool tripods. There are two primary characteristics of these objects. Firstly, they often feature relatively thin supporting structures, which can present sophisticated difficulties for collision avoidance. Secondly, they typically exhibit inconsistent shapes across various heights. Thus, addressing these objects with specific caution regarding both collision avoidance and robot sensing is essential.

A significant concern in recognizing and categorizing these four types of construction site objects and materials lies in the intricacies of robot sensing, particularly regarding how these elements may affect robot navigation. Recent advancements in field robotics frequently incorporate LiDAR sensors as part of the sensory apparatus for collision avoidance. The appearance of construction objects to different types of LiDARs, or their appearance as seen from various LiDAR placements, especially at the robot's height, can significantly influence the robot's navigational performance. Consider the example of open twin-step ladders (Figure 3.3), which typically assume a triangular shape. When utilizing a 2D LiDAR, different horizontal scans will be generated at different heights. If the LiDAR is positioned at the robot's base, close to the floor, it might scan an unoccupied space between the ladders that appears wide enough for the robot to pass through. However,

such a scenario is highly perilous and must be diligently avoided. Therefore, it is imperative to design a sensory package that is capable of accurately capturing and interpreting the essential features of the site objects.

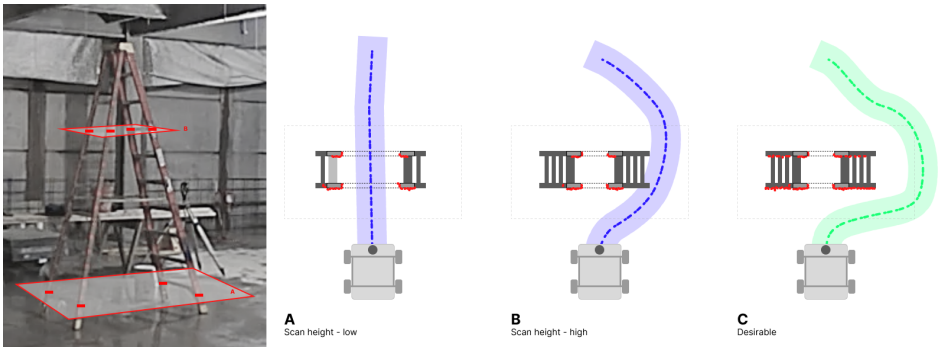

Figure 3.3: Case of open twin ladders: (A) scanning at a low height risks the robot passing underneath the ladder; (B) scanning at a high height risks the robot being tripped by the ladder; (C) shows the desirable 2D scan representation for collision-free navigation.

3.2.2 Carpentry Formwork Tools and Hardware

Based on camera images and dimensional measurements (Figure 3.4) collected on-site by my colleague Emek Erdolu, the specific tools and materials utilized in carpentry formwork are demystified and elucidated. In this extended analysis, I additionally sourced data from the internet, i.e. product catalogs on Amazon, Home Depot, Lowes, etc, to validate the approximate weights of each type of tool and hardware/material (Table 3.1). Delving into these items in greater detail serves not only to deepen the understanding of carpentry formwork activities but also to gauge the preliminary implications for the robotic framework, particularly concerning its load capacity, spatial requirements, and potential limitations.

The materials and hardware frequently moved between the work zone and the central workbench comprise (1) boxes of bolts for anchoring formwork panels, (2) long metal braces designed to support concrete forms, (3)

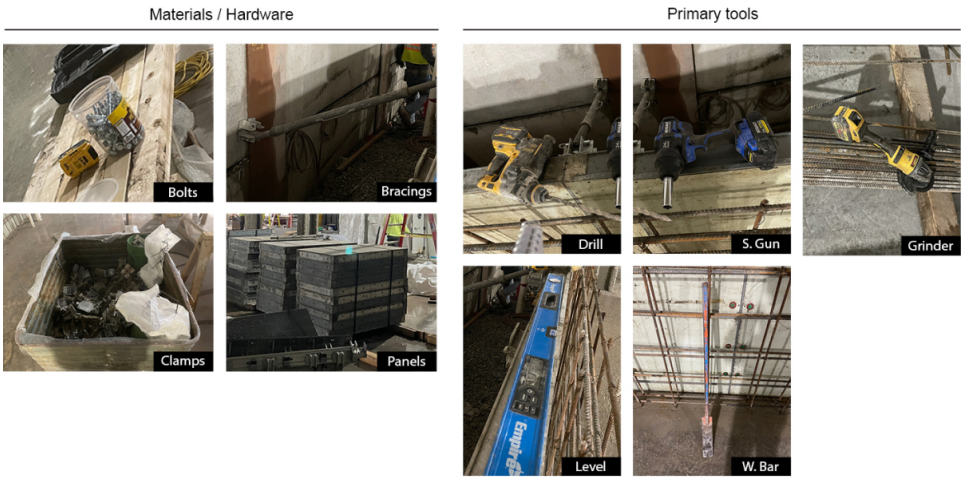

Figure 3.4: Site photos of tools and hardware in carpentry formwork. Image credit: ReAC Research Group, 2022, photos by Emek Erdolu.

metal clamps for quickly securing and tying formworks together, and (4) aluminum slab panels. In conjunction with the installation of these materials and hardware, the tools transported include heavy-duty items such as (1) impact wrenches for fastening bolts, (2) rotary hammer drills (with and without extra-long drill bits) for creating anchoring holes, (3) magnetic box levels (both long and short) for checking and marking references, and (4) claw end wrecking bars for removing nails and lifting. In addition to these heavier tools, there are also more common, smaller tools like pliers, hammers, and tape measures, which are essential for hand detailing and fine adjustments.

In the observed convention of manual transport, smaller items like bolts and bracings are often carried by hand, while heavier, longer, and bulkier items necessitate the use of trailing carts or other facilitations for mass transportation. Building upon this practice, I discussed with the research group to pinpoint items that would feasibly be transported by the envisioned robot support. These items encompass bolts, clamps, impact wrenches, rotary hammer drills, pliers, hammers, and tape measures. Items that cur-

Category	**Item**	**Approx. Size** (cm)		**Approx. Weight** (kg)
		L	**W**	
Hardware/	Box of bolts	24	12	1.5∼2.5
Material	Bracing	110∼120	–	3.5∼4.5
	Rapid Concrete Forming Clamps	10∼11	7.5	0.3∼0.5
	Aluminum Slab Panel	110∼120	40∼60	22∼25 $/m^2$
Tool	Cordless Impact Wrench (Kobalt 24V)	21	16	1.5
	Rotary Hammer Drill (DeWalt 20V)	30	25	2
	Rotary Hammer Drill (DeWalt 20V) w/ Long Drill Bit Attached	70	25	2.2
	Magnetic Box Level (Short)	61	7	0.8
	Magnetic Box Level (Long)	121	7	1.4
	Claw End Wrecking Bar	116	9	1.2∼2.5
	Plier	24	7	0.3∼0.4
	Hammer	36	11	0.4∼0.6
	Tape Measure	8	7	0.3

Table 3.1: Size and weight estimate of carpentry formwork tools and hardware.

rently require trailing carts for bulk transport, such as aluminum slab panels and bundles of bracings, are not considered within the scope of the current research due to the added complexity they would introduce in terms of navigation and load carriage. Nonetheless, exploring the development of trailing capabilities could be a worthwhile avenue for future research to enhance the practical, weight-bearing capabilities for robot navigation.

Additionally, slim, lightweight, and long items such as levels and wrecking bars are temporarily excluded from the scope of this research. The reason for this exclusion is that the standard dimensions of most Unmanned Ground Vehicle (UGV) robots are not designed to accommodate items that are thin and long. While the robot could technically carry these items based on weight, their safe transportation would likely require a specially designed, vertically oriented bucket. Placing these items horizontally on a navigating robot without a secure fastening mechanism could lead to the items protruding beyond the robot's profile, risking them falling off or potentially causing collisions.

Given these initial considerations and the identified scope, as well as the common size and weight of the analyzed tools and hardware, the envisioned carriage space for the robot is anticipated to be a cubic area measuring approximately 35∼45 cm in length, positioned on a flat surface atop the robot. To provide meaningful support, this space should be capable of accommodating a medium pile of clamps, one to two boxes of bolts (40∼50 per box), and two to three hand tools, with an overall maximum weight capacity of about 10∼15 kg. With this capacity, a single trip by the robot could approximately supply materials for 10∼20 minutes, based on a rough estimate drawn from observing the time-lapse recordings. Should there be any insufficiencies in the anticipated carriage capability, I would be addressing the current limitations through more frequent and adaptive robot runs, rather

than bluntly increasing the carriage size.

With these rough estimates in mind, a case that warrants specific attention is the rotary hammer drill, particularly when it is equipped with an extra-long drill bit. This configuration can significantly increase the tool's length to approximately 70cm, leading to the aforementioned protrusion issues. Consequently, a practical consensus among the research group involves the advisory of detaching the drill bit, which measures around 35cm, from the drill body, which has dimensions of 30cm by 25cm. Once separated, these parts are anticipated to comfortably fit within the proposed robot carriage space, thus meeting spatial limitations and facilitating safe transport.

3.3 Specificity of Formwork Activities

At first glance, carpentry formwork activities appear chaotic and complex. However, through careful observation of raw time-lapse camera recordings and qualitative insights from team interviews and conversations with workers, I discern a method to decompose these activities from both individual and collaborative standpoints. This decomposition does not oversimplify or diminish the intricate tapestry of carpentry formwork activities but rather explores them through the technical lens of envisioned robotic support. Specifically, I identify three distinct individual work paradigms and two collaborative work paradigms. These insights serve as crucial foundations for developing the robotic framework, training the contextualized social navigation algorithm, and designing worker-robot interactive experiments and evaluations.

3.3.1 Individual Work Paradigms

In social robot navigation, pedestrian behavior is often assumed to be goal-oriented and ambulatory, with agents moving efficiently toward their destinations. However, within labor-intensive work environments, worker movement patterns exhibit significantly greater variability and are closely linked to the tasks being performed. This variability necessitates a rethinking of traditional agent modeling approaches to accommodate the unique demands and characteristics of construction workers. These insights underscore the importance of carefully abstracting and decomposing worker activities, which serves as an essential foundation for developing effective navigation algorithms.

The proposed classification of individual worker activities into three paradigms (Figure 3.5) provides a framework for understanding and modeling the distinctive nature of carpentry formwork tasks. These paradigms bridge the gap between the specific characteristics of formwork installation workers and the more general representations used in robotics research, such as the Bird's Eye View (BEV) for path planning and motion prediction. In BEV, workers——referred to as agents——are modeled as circles with three key navigation attributes: safety radius, position, and velocity. By situating formwork activities within these paradigms, we can draw meaningful comparisons to existing pedestrian-focused robotic navigation scenarios and highlight the unique challenges presented by the observed construction tasks.

- **Paradigm I - Work Adjacent to the Formwork Wall**: This paradigm encompasses manual activities performed near the target wall, typically involving both Worker A and Worker B. Specifically, it includes instances such as (1) Worker A installing and reinforcing aluminum panels on the rebar cage, (2) Worker B preparing panels

against the wall, and (3) Worker B handing tools and hardware to Worker A. This paradigm is characterized by significant variations in the worker's safety/work radius, which must also account for the objects being handled to prevent collisions with the robot. The movement patterns in this paradigm are slow-paced and mostly linear, occurring adjacent to the wall, with potential stop-and-go behavior reflecting varying velocities as observed from a BEV.

- **Paradigm II - Walking/Traversing the Site**: This paradigm pertains primarily to Worker B's movements, including (1) transporting tools and materials from the workbench to the current work zone, and (2) moving panels, tools, and hardware between work zones. It also encompasses the migration of workers to subsequent work zones. Although these movement patterns may resemble typical pedestrian behavior, the worker's safety radius needs to be larger to accommodate the carried objects.

- **Paradigm III - Installing in a Focused Area**: This paradigm applies to tasks that require sustained or dedicated installation efforts, leading Worker A to remain nearly stationary for extended periods, along with moments when Worker B temporarily stands by while preparing panels or materials. It is characterized by minimal, fluctuating movements within a confined area, near-zero velocity, and a varying safety radius. These features highlight scenarios where worker mobility is significantly reduced, reflecting the focused and stationary nature of carpentry formwork installation activities.

Figure 3.5: The three individual work paradigms: (1) work and progress by the formwork wall, (2) traversing the site, and (3) installing and standing still.

3.3.2 Collaborative Work Paradigms

Given the collaborative nature of carpentry formwork, it is also essential to consider the specific dynamics of the worker duo's teamwork in the development of interactive robotic support. This consideration is particularly crucial when examined alongside the previously outlined three individual work paradigms and the characteristics of the worksite layout. As a result, two collaborative paradigms are identified: interaction and separation (Figure 3.6).

- **Paradigm I - Interaction**: This scenario primarily unfolds near the target wall within the current work zone, where Worker B is actively passing panels, tools, and materials to Worker A. Worker A is predominantly engaged in tasks adjacent to the wall (Individual Paradigm I and III), while Worker B may find himself in any of the three individ-

ual paradigms: working (Individual Paradigm I), moving (Individual Paradigm II), or waiting (Individual Paradigm III). Typically, the distance between the two workers does not exceed 3m. As mentioned earlier, the site layout usually features the work zone encircled by material bulks, with both workers operating in a corridor-like space approximately 2.5 to 3 meters wide, parallel to the target wall. Robotic support in this scenario must be designed to adeptly function within a tightly constrained area, facilitating close interactions between the workers without adding to the congestion.

- **Paradigm II - Separation**: This case arises when Worker A and Worker B function independently on the work site, temporarily without the need for direct collaboration or simultaneous task execution. For instance, Worker A may focus on installing panels, while Worker B exits the current work zone to fetch or prepare for aluminum panels. This scenario underscores the necessity for a robot that can efficiently navigate the workspace, assisting without obstructing the activities of either worker. This paradigm presents a relatively simpler technical challenge compared to the previous one.

The collaborative paradigms underscore the necessity of developing a robotic framework that is both aware of the physical layout of the worksite and sensitive to the complex patterns and dynamics of human teamwork. They also provide a contextual basis for designing algorithmic experimental cases and physical evaluation scenarios (see Section 4.5 in Chapter 4), serving as important criteria for developing and examining the proposed robotic support.

Clarification of these paradigms also highlights the unique aspects of worker density that diverge markedly from the focus of most current social navigation research. Typically, the area of interest is primarily occupied

Figure 3.6: The two collaborative work paradigms: (1) interaction and (2) separation.

by two workers, demonstrating heterogeneous activity patterns in notably congested workspaces. On occasion, an additional one to two workers from other construction trades might enter the area to obtain tools or simply to pass through. Consequently, the maximum number of individuals anticipated within the area usually does not exceed five. This feature hints at a research focus apart from prevailing studies, which typically assume a homogeneous dynamic among a dense crowd of 15 or more pedestrians. It emphasizes a shift towards exploring interactions within a smaller, more complex, and diverse group of individuals, highlighting the nuanced and intricate nature of social navigation within such specialized settings.

3.4 Context-Specific Technical Challenges

Besides the aforementioned material- and social-specific challenges, the above analysis also leads to several high-level technical concerns. The foremost concern pertains to safety. Given the human-centric positioning, the robotic support is anticipated to function within close quarters of carpentry workers,

blending into their established workflow. Such expected proximity underlines the critical need for prioritizing safety as a fundamental prerequisite to implementing any supportive functionalities. Instead of seeing safety only as an issue for collision avoidance algorithms, it is essential to devise multi-tiered solutions that harmonize both hardware and software efforts. At times, it may be necessary to implement compensating measures to ensure a layer of cascaded redundancy. Furthermore, establishing an iterative, case-specific approach for the swift enhancement of the system is crucial, ensuring immediate responses to new issues identified through both virtual and physical testing.

The second concern is closely tied to developing a technical roadmap that is both effective and efficient for the specific carpentry formwork context. Given constraints such as limited research time, manpower, and funding, it is essential to identify practical solutions that are tailored to the unique requirements of this environment, avoiding the deployment of excessive technical capabilities that do not offer direct benefits. A clear example of this principle is the selection of the navigation stack. In the carpentry formwork site, the floor is typically flat, as concrete leveling has already been completed, indicating that most navigation challenges can be efficiently addressed within a 2D navigation space. This contrasts with more spatially adaptable platforms like drones or quadruped robots, which target 3D navigation environments. However, indoor construction sites still contain geometrically complex 3D objects, and the floor may not be as clean or smooth as in environments where 2D navigation has been widely studied. Thus, effectively navigating this space——while incorporating an appropriate level of 3D awareness and retaining the computational efficiency of a 2D navigation system——requires a carefully designed engineering architecture that leverages suitable tools, methods, and resources.

The third challenge is in the specificity of short-range navigation granularities. Unlike autonomous vehicles or drones that navigate expansive outdoor terrains, the robot support system envisioned is designed for indoor, confined, populated spaces, with operational distances expected to be less than 30 meters. This distinct operational context steers the research focus from capabilities suited for long distances, such as extensive mapping, large-scale location semantics, or GPS-based localization, towards emphasizing short-range precision and adaptability. The development of the robot framework will prioritize equipping it with sensing and planning capabilities finely attuned to the immediate environment. This includes being responsive to the dynamic presence of workers and site objects within a proximity of approximately 3 to 5 meters. Such short-range focus necessitates addressing specific issues of different themes: accurately distinguishing workers against a complex background, maintaining precise self-localization within a gradually evolving worksite, and adaptively reacting to unforeseen appearances of workers or objects. Addressing these issues would require both system-level and component-level handling.

These challenges provide a preliminary direction and establish fundamental guidelines for the technical research to be conducted. Consequently, these context-specific considerations will be continually referenced throughout the subsequent chapters. In the next chapter, I will delve into the specific methods employed and utilized in this research, aiming to partially or fully address the raised concerns. This will include an exploration of both the hardware system and the software architecture developed to navigate the complexities of the research context.

Chapter 4

Methods

At a high level, the research adopts a research-by-prototyping method [Budde et al., 1990, Yan and Gu, 1996], moving beyond theoretical description to outline a concrete technical framework. A minimal viable robot prototype is developed as a physical realization of this framework. Conceptually, the prototyping process is structured into three main phases: (1) technical design and method selection, (2) major iterations and evolutions, and (3) minor optimizations and finalizations. These phases are integrated into the iterative research process, shaping both the system architecture and the development of technical components. This chapter not only outlines the key decisions made regarding different technical aspects but also documents significant trials and iterations that shaped the final system, providing a comprehensive view of the prototyping trajectory.

From the outset, reinforcement learning (RL) is chosen as the primary technical backbone due to its potential for enabling socially aware and adaptive robot navigation around workers on-site [Chen et al., 2017, Everett et al., 2018]. This decision, made in tandem with ongoing discussions within the ReAC research group regarding the specifics of robotic support scenarios

and functions, is not based on intuition alone. This decision stems from three key considerations.

Firstly, as emphasized in Chapter 2, RL is well-suited for environments characterized by stochasticity, uncertainty, and complexity—hallmarks of on-site indoor construction settings. Traditional methods, such as rule-based approaches, struggle to handle the intricate and dynamic nature of real-world worker activities, making RL a promising alternative. RL's capability to develop a decision-making policy through learning from trial-and-error experiences provides a significant advantage in such dynamic environments.

Secondly, there is a robust body of existing research and resources in RL for social robot navigation, offering a strong basis for this study. Notably, the CMU BIG research group provides an open-sourced simulation environment [Shum et al., 2022] that facilitates the generation of large datasets required for RL training, addressing a common bottleneck in RL-based approaches. Additionally, a preliminary review of the literature reveals extensive work on social robot navigation in other contexts, particularly in dense crowd settings, which, while different from construction environments, offers valuable foundation and progress in the field.

Lastly, drawing on insights from our multi-drone research [Fang et al., 2020], we recognize the potential to incrementally fine-tune and adapt an RL model through multi-stage curriculum training. This approach offers a crucial advantage by enabling the robot's generic navigation behavior to be more closely aligned with worker activities and site-specific features, contrasting with conventional methods that often provide limited flexibility for addressing these intricacies, aside from the adjustment of a few hyperparameters.

With this pivotal decision made, the envisioned reinforcement learning setup requires the following state information for real-time decision-making:

(1) the robot's position and velocity, (2) the position, velocity, and radius of all nearby agents (i.e., workers), (3) the relative position of the intended goal, and (4) the robot's heading angle and preferred speed. Notably, while these states can be easily acquired in simulation environments, real-world robots must derive this information from raw sensory data. Addressing this challenge requires not only algorithmic efforts but also a sophisticated systems engineering approach to the robotic framework, which is central to the research-by-prototyping methodology.

Mathematically, the problem is framed as a Markov Decision Process (MDP) $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \mathcal{S}_0 \rangle$, drawing from insights in social robot navigation research [Chen et al., 2017, 2019].

More specifically, the joint state observable to the robot consists of two parts:

$$\mathbf{s}_t^{\text{joint}} = [\mathbf{s}_t^{\text{full}}, \mathbf{w}_t^{\text{obs}}].$$

The first part, $\mathbf{s}_t^{\text{full}}$, at time t, includes the robot's position $\mathbf{p} = [p_x, p_y]$, velocity $\mathbf{v} = [v_x, v_y]$, goal position $\mathbf{p}_g = [p_{gx}, p_{gy}]$, radius r, heading angle θ, and preferred speed v_{pref}. The second part, $\mathbf{w}_t^{\text{obs}}$, includes the positions $\mathbf{p}_i = [p_{ix}, p_{iy}]$, velocities $\mathbf{v}_i = [v_{ix}, v_{iy}]$, and radii r_i of all nearby agents (i.e., workers)[1].

The action space, $\mathcal{A}$, consists of the robot's possible movements $\mathbf{a} = [a_x, a_y]$, with the assumption that $\mathbf{a}_t = \mathbf{v}_t$. The transition function, $\mathcal{P}$, governs the environment's dynamics, determining the next state based on the current state and action. The reward function, $\mathcal{R}$, provides feedback to the robot based on its actions, encouraging socially comfortable navigation around workers while avoiding collisions. The discount factor, γ, balances

[1] An important distinction to note here is that the robot is unaware of the worker's intended goal. This aligns with the common assumption in social navigation research, where agents can only know their own intent unless explicitly communicated.

immediate and future rewards, and $\mathcal{S}_0$ denotes the initial state of the robot.

This research aims to solve the MDP using reinforcement learning, optimizing the expected return:

$$R_t = \mathbb{E}\left[\sum_{k=t}^{T} \gamma^{k-t} r_k\right].$$

With the context analysed in Chapter 3, the objective expands to optimize this return while considering the cluttered site layout L and unique worker behavioral patterns W.

Under this formulation, and in parallel with evolving qualitative discussions within the group—especially concerning the specificity of carpentry formwork as a context for tool/material-delivery robotic support—an initial modular robotic framework is established to lay out the technical roadmap (Figure 4.1). The modular design serves three key purposes: (1) orchestrating a pipeline for real-time execution of RL on the robot by processing and supplying the required input state from raw sensory data, (2) complementing the current RL-based social navigation algorithm, particularly its limitations in navigating geometrically complex construction environments, and (3) supporting flexible, module-wise adjustments and improvements without structurally impacting the overall system. In more detail, the framework consists of five key modules:

1. **Robot Hardware and System Setup**: This module involves the creation of a physical robot prototype, tailored to the specifics of carpentry formwork, incorporating the necessary hardware components (such as sensors, processors, and actuators) to support real-time execution of the RL algorithm and interaction with the construction workers and environment.

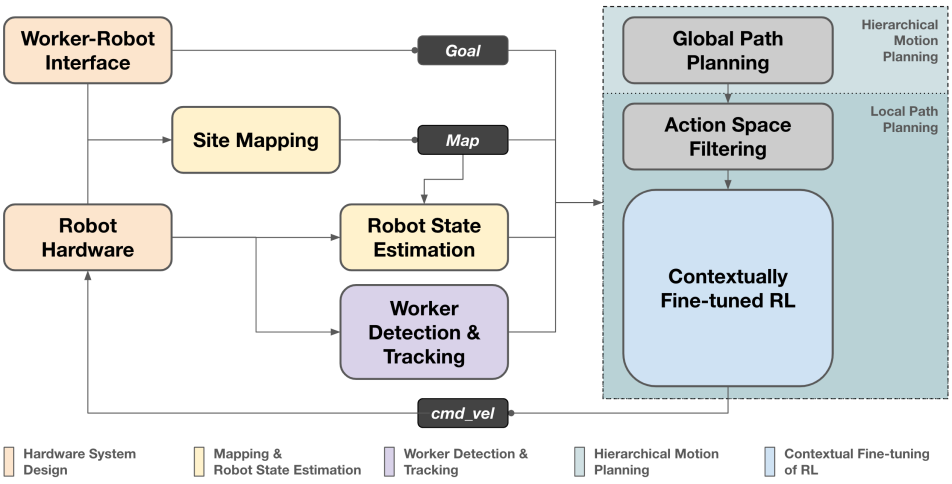

Figure 4.1: System design of the RL-driven robotic framework [Wu et al., 2024b].

2. **Unstructured Site Mapping and Robot State Estimation**: This module focuses on enabling the robot to accurately perceive and interpret the complex, dynamic construction site environment. It includes the development of algorithms and systems to gather and process raw sensory data, thereby accurately estimating the robot's state (position, velocity, orientation) in relation to the surrounding environment.

3. **Worker Detection and Tracking**: This aspect involves developing robust methods for identifying and continuously tracking the movements of carpentry workers within the robot's operational vicinity. It is critical for ensuring the robot's socially aware navigation and safety in the dynamic construction site, where workers' movements are varied and often unpredictable.

4. **Hierarchical and Efficient Motion Planning**: This module aims to develop a hierarchical structure that efficiently integrates classical collision avoidance methods with RL-based social navigation methods, for the purpose of a robust pipeline for both safe and social navigation in complex and populated construction environments.

5. **Contextualized Finetuning of RL Model**: This involves adapting and optimizing the reinforcement learning model to the specific challenges and nuances of the carpentry formwork context. It includes adjusting the RL algorithm to better understand and respond to the unique patterns of worker movement and the spatial complexities of the construction site, ensuring socially comfortable worker-robot interactions remain contextually effective.

To ensure integration of the modules within the system architecture, the Robot Operating System (ROS) [Quigley et al., 2009] —a widely recognized open-source software framework—is selected at the outset of this study to ensure low-latency communication between module inputs and outputs. Particularly, the ROS navigation stack, a technical stack frequently utilized in the development of various 2D navigation robots, is exploited as an initial skeleton of the research. This stack provides a basic pipeline that utilizes raw sensor data for mapping, localization, and navigation tasks. However, for the purposes of this research, this navigation stack is overly simplistic and inadequate. The modules described above build upon these fundamental functionalities by introducing critical layers of customization essential for the development of the envisioned robotic support.

This chapter proceeds by offering a detailed description of each module, including the selection of technical approaches, the specificity of engineering details, and the considerations in addressing context-specific challenges mentioned in Chapter 3. Finally, this chapter concludes with a summarizing system dashboard that details the full engineering details of the entire framework, showcasing how modules are integrated to work towards the envisioned support, and how a conceptual framework is consolidated into a sophisticated working robot system.

4.1 Robot Hardware Design

As the tangible foundation of the research, the development of the robot's hardware system undergoes three primary stages of iterative refinement (Figure 4.2). This evolutionary process mirrors the decisions and adjustments made on the software side, encapsulating a progressively meticulous synthesis.

- **Stage 1 in Figure 4.2, Preliminary Setup:** Initially, the robot's base and onboard computing environment are configured to support basic control and operation (e.g., teleoperation using keyboards and joysticks). Two LiDAR sensors are provisionally affixed to the robot base to facilitate physical experiments and pipeline walkthroughs of basic ROS navigation stack components, such as `gmapping` and `AMCL`, which require `laserscan` inputs.

- **Stage 2 in Figure 4.2, Rapid Prototyping:** This phase is dedicated to the rapid testing, integration, and iteration of hardware related to the aforementioned modules. It commences with an enhanced design of the robot's metal frame using T-slot construction, building upon the proof of concept from Stage 1. The frame is improved in terms of height, usability, and strength through a double-frame structure, thereby augmenting stability for accommodating various sensors. A significant addition is the establishment of a horizontal metal beam plate designed to support carrying loads atop. During this stage, functionalities associated with 3D LiDAR are extensively refined, culminating in the upgraded design of the 3D LiDAR and IMU modular box. Furthermore, a depth camera is experimentally attached to the robot frame to allow for flexible testing of worker detection and tracking algorithms.

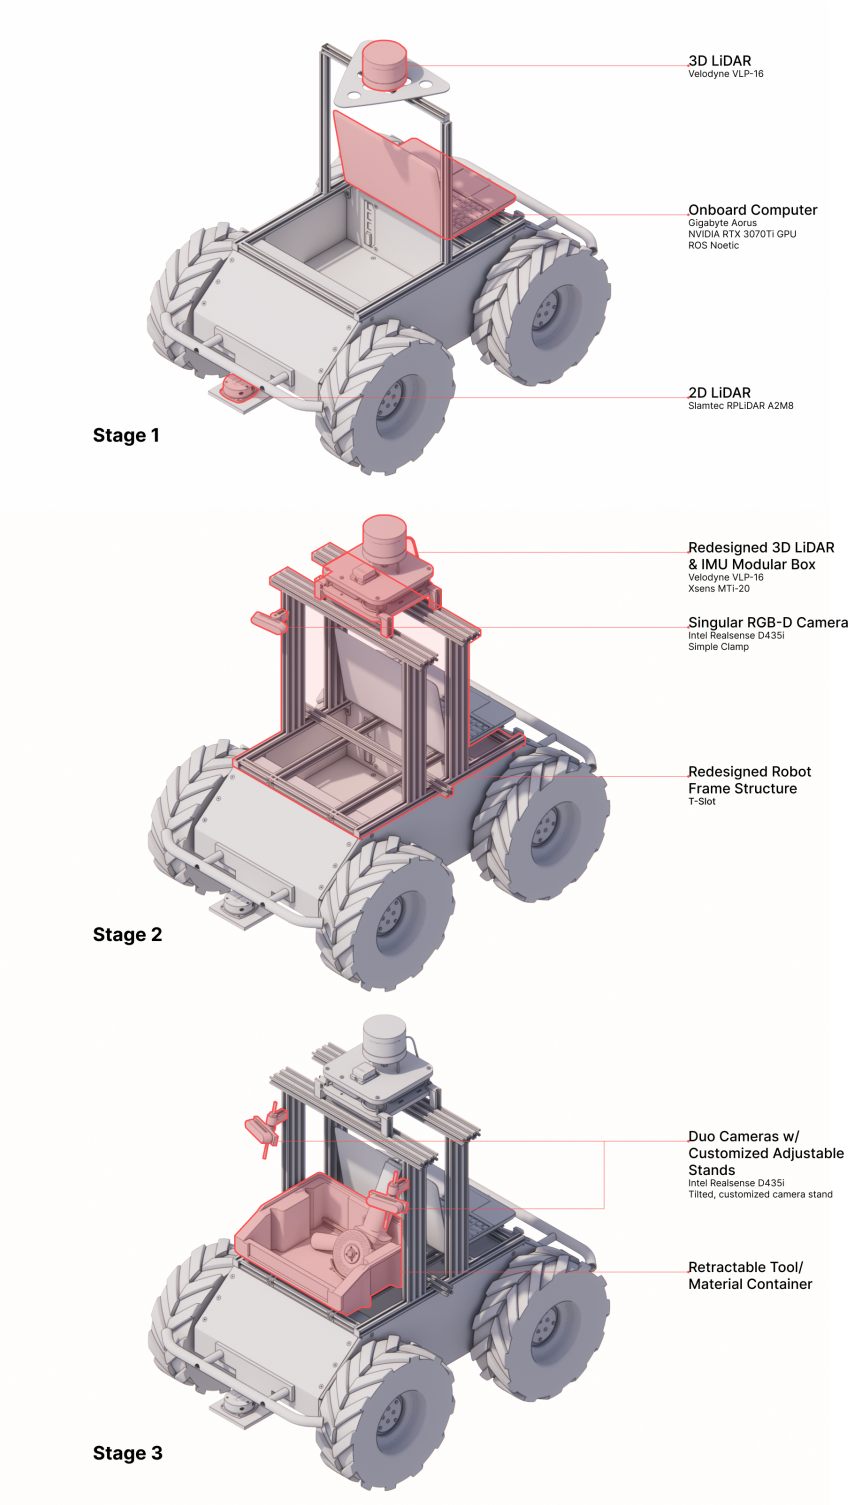

Figure 4.2: Three major stages in robot hardware design and development.

- **Stage 3 in Figure 4.2, Finalization:** With most software development completed, this stage sees the final optimization of the robot hardware building, including the introduction of two depth cameras as the final solution for worker detection and tracking. These cameras are incorporated using 3D-printed, custom-designed adjustable stands, firmly affixed to the vertical frames. Collaborative efforts within the ReAC research group also lead to the conceptualization, testing, and finalization of the tool/material container, taking into account both material and social considerations.

The following sections provide a more detailed exploration of specific components such as the robot base, LiDARs, cameras, etc., to elucidate the rationale behind such iterations and design choices.

4.1.1 Base and Chassis

Selecting an appropriate movement base/chassis involves a laborious experimental process among three different models (Figure 4.3). Initially, we utilize a Fetch robot, commonly known for its application in contexts such as factories and indoor logistics. Despite its integrated motherboard, 2D LiDAR, and depth camera, the Fetch robot offers limited compatibility for further customization. Further challenges, including malfunctioning LiDAR and motherboard issues, eventually lead to its discontinuation in the research.

Subsequently, a custom-built[2], triangular, omnidirectional robot base named Rocky is examined. Equipped with three mecanum wheels, Rocky

[2]The Rocky robot, developed by students from the Robotics Institute at Carnegie Mellon University, has been pivotal in the research conducted by the Robot Intelligence Group. It features a vertical, slim profile with a microphone at the top for voice interactions in human crowds. In this research, we primarily use the base part of the robot, repurposing it for navigation on construction sites.

offers quite flexible mobility. However, its control system lacks stability, and measures to limit the robot's maximum speed are not reliably effective. This often results in the robot moving too quickly, posing risks of uncontrolled drifting and collisions. During tests, there is a necessity for a person to intercept the robot within seconds, highlighting significant safety concerns. Therefore, Rocky is also deemed unsuitable for the context.

Fetch

Figure 4.3: Comparison among three robot bases.

The Clearpath Husky A200 is ultimately chosen, marking a significant milestone in the research. Originally designed for rugged outdoor environments, the Husky robot base excels in adaptability to context-specific challenges, distinguished by three key attributes: physical suitability, customization potential, and safety. Notably, these attributes can also serve as useful baseline criteria for other research endeavors seeking appropriate robot bases for customizing robots in human-centered contexts.

1. **Physical Suitability**: Husky's sturdy design, featuring four durable, thick rubber wheels, is ideally suited for navigating the challenging conditions of construction floors, which frequently include obstacles like dents, cables, and debris, as analyzed in Chapter 3. Its wheels are designed to resist hazardous materials and conditions, such as water puddles mixed with dust and concrete powder, nails, and sharp metal fragments. Additionally, Husky's solid mid-sized build and low

center of gravity make it highly suitable for transporting carpentry tools and hardware across slightly bumpy surfaces. Its width, comparable to or slightly wider than an average person, enables adequate maneuverability through narrow passages among construction objects and layouts, offering an advantage over larger, more physically capable UGVs. The inclusion of T-slot rails on the robot's rim also simplifies the attachment of additional structures atop through simple fixtures.

2. **Room for Customization**: Husky offers substantial flexibility for experimental customization. Its control and navigation stack, based on ROS, facilitates easy modifications and aligns with our framework's setup. The open-source nature of this stack and its compatibility with most Ubuntu laptops enable extensive customization of the robot's onboard computer. Moreover, while the obtained Husky robot base initially lacks a comprehensive sensor suite, it provides a straightforward platform for adding various sensors, both in terms of physical connection and software configuration.

3. **Safety**: Husky's design incorporates multiple features that help address the safety concerns discussed in Chapter 3. It includes an easily accessible emergency stop button, extensively used throughout development, and a simple method for limiting maximum speed via hyperparameter adjustments. These features ensure quick intervention in potential collision scenarios, allowing immediate cessation of the robot's movement by pressing the emergency stop.

The Husky robot base is then equipped with a laptop featuring an NVIDIA RTX 3070 GPU as its main onboard computer (see Stage 1 in Figure 4.2). This configuration is intended to prepare for the acceleration of real-time Deep Neural Network (DNN) inference, crucial for modules such as worker detection and tracking and reinforcement learning.

4.1.2 LiDARs

Drawing inspiration from relevant domains such as autonomous driving and indoor social robotics, Light Detection and Ranging (LiDAR) units are identified as essential early in the research. Given the unstructured nature of construction sites and their cluttered environments, LiDAR sensors are invaluable for their ability to rapidly detect nearby objects and people while accurately measuring their relative positions. This capability is critical for tasks such as site mapping, robot localization, and path planning. Motivated by their successful application in autonomous driving, the research initially adopts a standard setup with a single 3D LiDAR, a Velodyne VLP-16, positioned approximately 0.8 meters above the robot on the metal frame (Stage 1), aiming for 360-degree environmental coverage.

However, preliminary lab testing reveals inadequacies in this configuration. The robot exhibits a significant likelihood of failing to detect low-lying objects within close proximity (<1m), as illustrated in Figure 4.4. A detailed investigation into this issue reveals a "blind cone" effect, attributable to the limited Field of View (FOV) caused by the bottom metal rim of the selected LiDAR model. This limitation is particularly problematic in our context, as highlighted by the "short-range navigation granularities" concern discussed in Chapter 3. Unlike other mid-range or long-range application scenarios where encounters with randomly placed, low-lying obstacles are less common or less problematic, our context regularly includes temporary, ground-level hindrances such as hand tools, material piles, and work residues (refer to Site Objects and Conditions in Chapter 3). Recognizing this as a non-trivial matter, it becomes imperative to devise a solution.

In response, intensive brainstorming within the research group leads to the decision to augment the initial setup by installing a 2D LiDAR, a Slamtec RPLiDAR A2M8, at the front of the robot, positioned as close to the ground

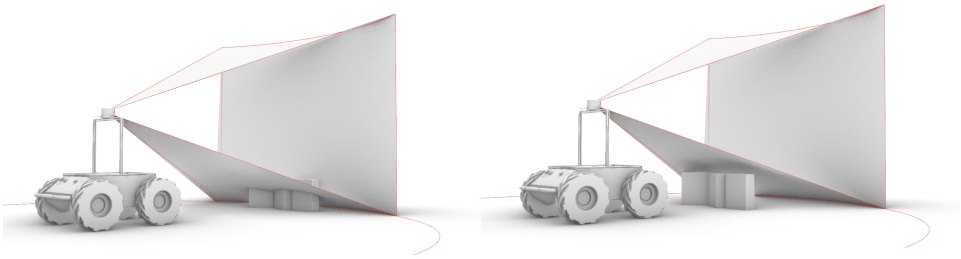

Figure 4.4: "Blind cone" effect caused by the metal rims of a 3D LiDAR. Image credit: ReAC Research Group, 2022, renderings by Jiaying Wei.

as feasibly possible, to complement the 3D LiDAR. Given the constraints of research resources, such as funding and effort, this adaptation is not intended to completely compensate for the "blind cone" in three dimensions. Instead, the 2D LiDAR aims to detect objects on a plane approximately 5cm above the ground. A geometric analysis of site objects (see Chapter 3) suggests this plane effectively scans the majority of low obstacles. For objects below 5cm, further examination of site imagery indicates that most can be safely traversed by the robot, including large construction debris, cables, and timber beams positioned on the ground for area delineation. Accordingly, the 2D LiDAR is mounted below the robot's front bumper on a custom-designed acrylic plate (Figure 4.2).

Besides addressing the "blind cone" issue, mapping and localization features also necessitate the tight coupling of a 9-axis Inertial Measurement Unit (IMU), an Xsens MTi-20, with the 3D LiDAR. This setup allows for the acquisition of corrective information concerning the robot's roll, pitch, and yaw. Drawing from insights in [Shan et al., 2020], this coupling leads to the development of a composite LiDAR box design (Figure 4.5). To achieve precise and orderly placement of the 3D LiDAR, its interface box, the IMU, a synchronization control board, and a collection of wires and cables, two robust acrylic plates are designed and fabricated. The box design not only

presents a neat profile but also introduces a pivotal functionality: its ability to be detached from the robot. This feature is particularly beneficial for conducting experiments related to LOAM (see in Site Mapping and Robot State Estimation Section), facilitating significant, independent testing sessions somewhat separate from the comprehensive robot experiments.

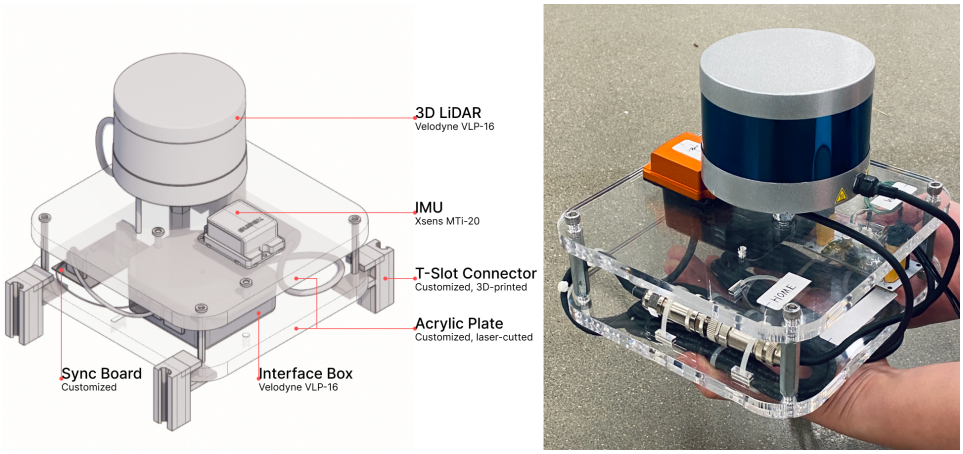

Figure 4.5: Composite LiDAR box design (left) and implementation (right).

4.1.3 Cameras

The assumption that "3D LiDAR will solve all issues" is reconsidered when determining the essential sensors for worker detection and tracking (refer to the Worker Detection and Tracking section). The high false positive rate in identifying workers against the complex backdrop of construction sites by utilizing a single 3D LiDAR necessitates the adoption of depth cameras. Unlike LiDARs or RGB cameras, depth cameras offer a cost-effective, intermediate solution, merging the benefits of laser scanning and visual recognition. However, a notable limitation, especially with the selected model, the Intel Realsense D435i, is its restricted Field of View (FOV). According to the official specifications, the horizontal FOV is about 85.2 degrees, and the vertical FOV is roughly 58 degrees.

This premise underpins the initiation of the prototyping phase, where one depth camera is initially temporarily mounted to the left side of the robot (Stage 2 in Figure 4.2). The connection utilizes a commercially available camera stand/clamp, allowing for flexible adjustment of the camera's orientation both vertically and horizontally. This adaptability is crucial at this stage, enabling exploration of optimal physical placements for enhanced performance in conjunction with the detection algorithm. Additionally, this stage involves thorough testing of the basic algorithmic pipeline, encompassing data collection and processing, model improvement, inference acceleration, and deployment as preparation for subsequent prototype refinement.

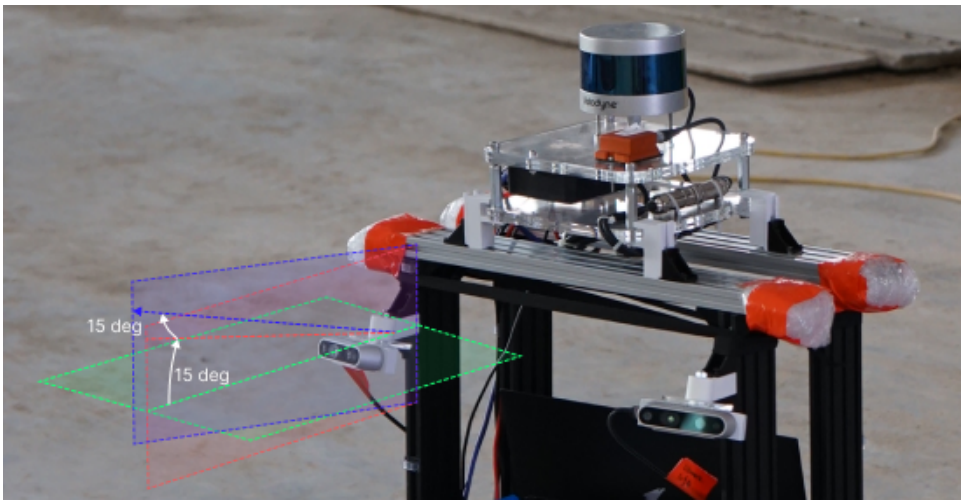

Figure 4.6: Tilted camera stand design and the duo-camera setup.

To address the FOV limitation inherent to using a single camera, two depth cameras are ultimately employed in tandem (Stage 3 in Figure 4.2). This setup, conceived in response to the "short-range navigation granularities" and "effective but excessive" concerns mentioned in Chapter 3, aims not to achieve complete 360-degree coverage around the robot but to cover, in good quality, even with some FOV overlap as redundancy, the robot's immediate operational area, i.e., its frontal space (Figure 4.6). In particular, the cameras are mounted on custom-designed, 3D-printed stands, ensuring

secure, stable attachment, crucial when navigating uneven construction surfaces that induce some degree of vibration. Based on empirical insights from earlier tests, the cameras are deliberately angled upwards by 15 degrees to improve the capture of workers' faces, especially when they are relatively close to the robot.

4.1.4 Additional Structure and Design

In addition to the hardware components critical for navigation, a collective effort within the ReAC group[3] leads to the design and implementation of a worker-friendly tool/hardware container and a Bluetooth communication module, complementing the full hardware design of the robot. Informed by analyses of the tools and hardware commonly used in carpentry formwork, as detailed in Chapter 3, discussions within the group resulted in a sturdy plastic container measuring approximately 35cm by 25cm, designed to be retractable over a 35cm by 45cm area on the robot's front deck's T-slot beam plate (Figure 4.7). This container aligns with the expectations set in Chapter 3, capable of holding 2-3 bags of metal hardware and at least 2-3 types of heavy hand tools, such as wrenches, rotary hammer drills, hammers, etc. The retractable feature considers both navigation stability and worker accessibility. For navigation, the container can be retracted, positioning its center of gravity within the robot's main body and beneath the upper metal structure for stability. During companionship phases, the container extends outward like a drawer, facilitating easy access to tools and hardware for workers, even when the robot maintains a certain safety distance from them.

To foster intuitive interaction with the robot, iterative brainstorming within the group led to the implementation of a keypad controller as the

[3]Jiaying Wei played a central role in these efforts.

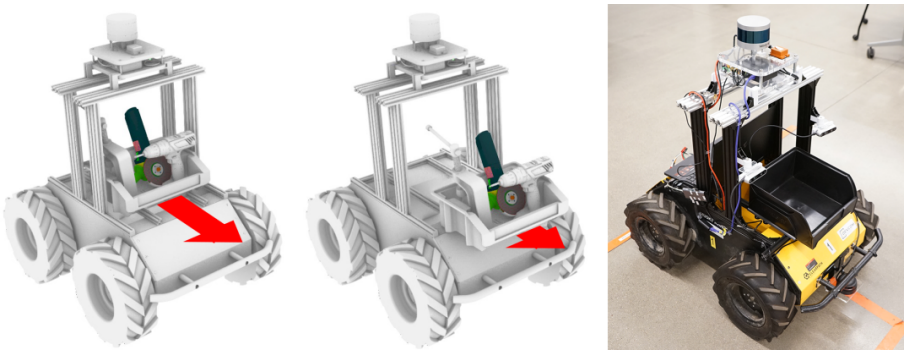

Figure 4.7: Retractable tool and hardware container. Image credit: ReAC Research Group, 2022, rendering by Jiaying Wei.

main Human-Robot Interaction (HRI) interface (Figure 4.8). While other interfaces and mediums, such as voice, gesture, and posture, are extensively discussed, they are ultimately not adopted. Construction sites, especially with the constant operation of noisy electric tools, are notably loud environments. Voice control, although seemingly convenient and intuitive, can be unreliable in application. The prospect of workers having to shout at the robot over the ambient construction noise is deemed undesirable. Similarly, the use of gestures or postures, despite their technical sophistication, would significantly disrupt carpentry workers' ongoing tasks. The most reliable and intuitive solution at the moment, as concluded from our discussions, remains the use of keypresses, which should be feasible even when workers are wearing their gloves. The Bluetooth controller designed by the group meets this requirement. This interaction method allows key locations within the studied context, namely work zones and the central workbench, to be efficiently memorized and encoded as numbers. Worker interactions with the robot are simplified to pressing the intended goal number followed by a confirmation key. Extensive lab testing confirms that the keys can be easily pressed even when wearing common work gloves.

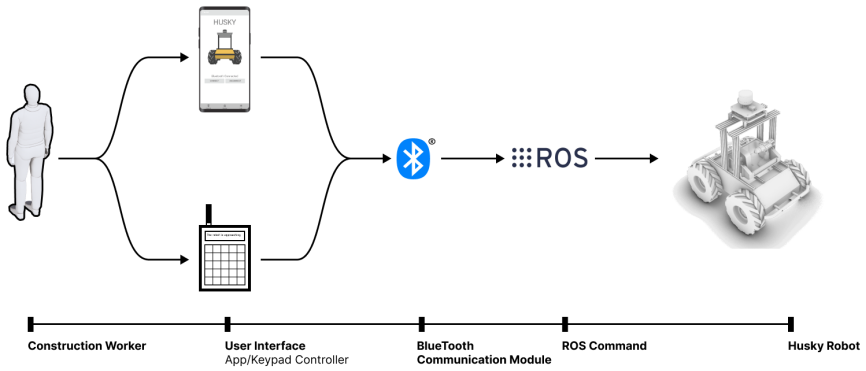

Figure 4.8: Custom-designed Bluetooth controller.

4.2 Unstructured Site Mapping and Robot State Estimation

The robot's effectiveness in navigating a dynamic site with shifting work zones depends on proficient mapping and accurate state estimation within the complex geometries of the construction environment. These efforts are geared towards resolving two principal questions. Firstly, how to aptly capture the site's complex 3D geometric features in a map that is useful for navigation without imposing excessive computational demands. Secondly, how the robot can utilize this map for effective self-localization during navigation. Highlighted in the "effective but not excessive" concern in Chapter 3, the support for carpentry workers is primarily designated for a flat work area, rendering the navigation task essentially a 2D path planning endeavor. Thus, tailoring the mapping and localization process to this level of simplicity and incorporating it within the prevailing 2D navigation framework and technical ecology is preferable. This approach circumvents the complexities associated with 3D navigation stacks with special focus on maneuvering through elevated and uneven terrains or engaging in 3D path planning—a feature that, while advanced, would be superfluous for our specific context

and could unnecessarily complicate the subsequent decisions and methods.

4.2.1 Mapping

The initial mapping attempt involves utilizing canonical 2D mapping tools commonly used in the ROS navigation stack, like `gmapping` [Grisetti et al., 2007] and `SLAM Toolbox` [Macenski and Jambrecic, 2021]. These packages use laser scans from 2D LiDAR to depict the site's layout as a 2D grid map, marking occupied spaces on a grid. In this iteration, site complexity is tentatively addressed by placing the 2D LiDAR as close to the ground as possible, under the assumption that unstructured site geometries primarily stem from objects irregularly placed on the ground. Capturing more features close to the ground is then thought to adequately represent the site's geometry. While this approach does resolve some challenges, it quickly proves overly simplistic and presumptive. Laboratory experiments reveal significant issues, such as the robot's tendency to navigate under high tables, posing a risk due to its inability to handle objects with height-varying shapes and profiles. In the grid map, large high tables are only captured as four dots, i.e., planar scans of four thin table legs, falsely suggesting enough space for the robot to pass through. Meanwhile, the slim table surface, at a height of approximately 0.9 to 1.2 meters, could sometimes be overlooked during navigation by the intervals of the 16 laser scan rings of the equipped 3D LiDAR[4]. This inadequacy risks mishandling objects like open ladders, tripods, slanted steel columns, and construction objects with extruding or height-varying features if actually applied on-site, as also mentioned in Chapter 3.

Acknowledging this shortfall, the second iteration shifts to representing the site as a 3D point cloud, addressing the previous limitations by directly

[4]The Velodyne VLP-16 LiDAR, as its model name suggests, uses an array of 16 infrared (IR) lasers to cover its 3D FoV.

capturing unstructured site objects in 3D. This approach employs LiDAR Odometry and Mapping (LOAM)-type methods [Zhang and Singh, 2014], utilizing 3D LiDAR to construct a 3D point cloud simply by teleoperating the robot around the area of interest for 1∼2 loop(s). These methods provide both real-time odometry estimates for robot state estimation and a reconstructed point cloud map of the site, a desirable by-product. In our context, the quality of the point cloud map, particularly its density, is critical to capturing the geometric details of construction objects. After testing various versions of LOAM methods, including the original LOAM [Zhang and Singh, 2014], LeGO-LOAM [Shan and Englot, 2018], and LIO-SAM [Shan et al., 2020], LIO-SAM is eventually chosen for its superior mapping quality, despite requiring a more complex hardware setup[5]. While the reconstructed point cloud features enough geometric detail of the site, how to utilize this rich information in tandem with subsequent path planning methods presents new challenges. Surveyed path planning algorithms and packages directly referencing 3D point clouds are either overly complex and redundant for our needs, or incompatible with previously tested framework modules, posing indispensable engineering difficulty. Additionally, the 3D navigation ecosystem in ROS is less mature than its 2D counterpart. These practical concerns necessitate an iteration of further optimization.

The third iteration draws insights from the previous two iterations by combining the simplicity of 2D grid maps with the detail of 3D point clouds. Using tools like `Octomap` [Hornung et al., 2013], points within a certain height range are flattened and merged into a 2D grid map. In more detail, the LIO-SAM generated point cloud is first preprocessed to remove points belonging to the ceiling and floor[6] before integration into a grid map. Laboratory

[5]As also mentioned in the "Hardware System Design" section, this method requires a tightly coupled 9-DoF IMU for essential initialization and corrective information

[6]If ceiling and floor points are not removed, the merged grid map would falsely mark navigable spaces as occupied.

experiments show that this map can effectively capture features like high tables and extruded structures in an appropriate fashion, while integrating with the framework's other tested modules, offering an ideal and balanced solution between 2D computational efficiency and 3D geometric richness, desirable for our context.

4.2.2 Robot State Estimation

The complex geometric conditions of construction sites also pose a significant challenge for robot state estimation. Accurately determining the robot's location within the map is crucial for planning its path to the designated goal, marked on the same map. Incorrect localization can lead to problematic robot behavior, such as colliding with walls or engaging in futile recovery actions like in-place rotation, posing safety risks to nearby workers if not properly addressed. Similar to mapping, the initial attempt at solving this challenge involves leveraging existing methods and packages in the ROS navigation stack, notably the Adaptive Monte Carlo Localization (AMCL) [Pfaff et al., 2006] approach. This method matches the 2D laser scans obtained during navigation with the pre-built grid map to infer the robot's most probable position. While effective in simpler environments, the conditions at construction sites introduce unique challenges due to their dynamic, non-stationary nature, with materials and objects frequently being moved around and changing shape as work progresses, rendering scan-map-matching methods less reliable. Laboratory testing reveals that, although the AMCL method generally performs well, changes in the placement of construction objects and materials such as boxes and blocks can sometimes cause localization drifts.

In response, the second iteration adopts a more comprehensive strategy that combines multiple estimations from different sources rather than

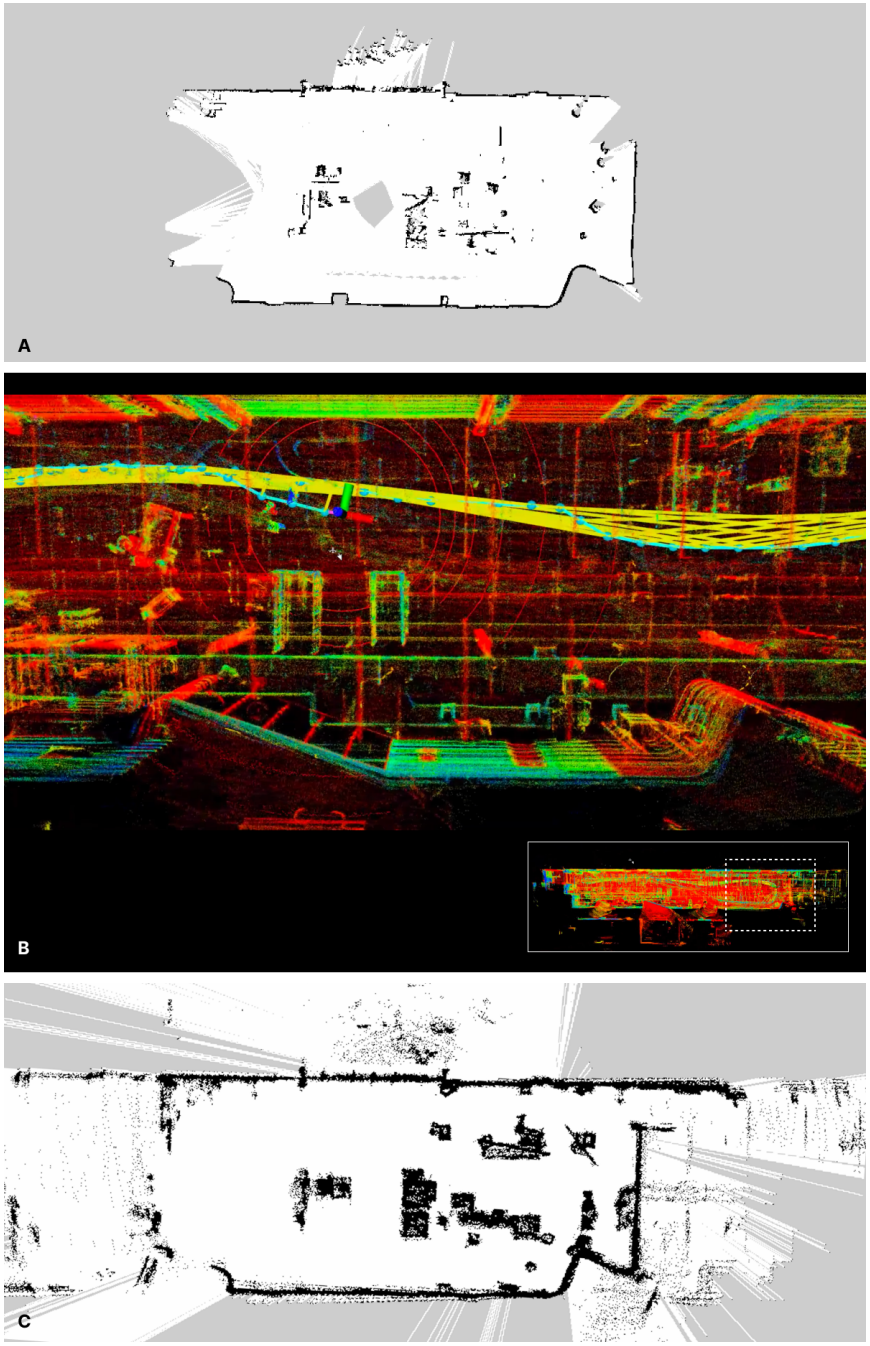

Figure 4.9: Three iterations of site mapping: (A) simple 2D mapping; (B) 3D mapping and point cloud reconstruction; and (C) processed 2D mapping from real-time point cloud reconstruction. Adapted from [Wu et al., 2024b].

relying on a single method. The high-quality odometry estimated by LIO-SAM serves as one such source, complemented by the odometry from the Husky robot's wheel encoders.[7] The odometry data from AMCL are also utilized, providing useful references in most scenarios. These diverse odometry sources are integrated using an Extended Kalman Filter (EKF) [Welch et al., 1995] to produce a consolidated estimate of the robot's odometry. Through empirical observation and comparison, the contribution of each odometry source to the EKF is fine-tuned, with the highest weight assigned to LIO-SAM odometry, followed by AMCL, and the lowest weight to Husky wheel odometry. Tested through laboratory experiments, once the robot's initial starting position is determined[8], this layered approach is relatively effective in practice. Nonetheless, localizing a robot within a dynamically changing construction site remains a complex challenge, presenting an intriguing area for future research.

4.3 Worker Detection and Tracking

As a critical precondition for planning the robot's next actions, it is crucial for the robot to rapidly identify surrounding workers amidst the complex milieu of construction elements. This challenge entails two key expectations: first, the anonymous identification of individual workers against the cluttered site backdrop, and second, the estimation of these workers' po-

[7]This source tends to experience slight drifting, and attempts within the group to improve its accuracy through hyperparameter tuning are not successful, making its use currently optional. If given more time and resources, however, this hyperparameter could potentially be fine-tuned by implementing a Proportional-Integral-Derivative (PID) controller.

[8]Determining the robot's initial position within the map is now performed by a research member offering a rough initial guess using the mouse. The implemented AMCL package then takes over and forms a more accurate estimate. This procedure in practice is quite intuitive, convenient, and effective.

sitions, velocities, and safety radii in relation to the robot[9]. Drawing on insights from previous research in social robot navigation [Everett, 2017], these expectations are addressed in two steps: detection and tracking.

4.3.1 Detection

In the detection phase, the primary goal is to recognize workers from raw sensory data and determine their relative state. Inspired by the success of 3D-LiDAR methods in pedestrian detection within autonomous driving scenarios, the initial research effort concentrates on LiDAR-based techniques. With the 3D LiDAR mounted atop the robot, renowned methods such as PV-RCNN [Shi et al., 2020] and PointPillars [Lang et al., 2019] are rigorously tested and implemented on the robot. The models are initially evaluated using well-known autonomous driving datasets like KITTI [Geiger et al., 2012], then deployed for lab experimentation. However, during laboratory testing, the models exhibit unsatisfactory performance, characterized by a high rate of false positive detections. Objects like irregularly stacked blocks, bare steel columns, vertical duct pipes, and wall fixtures are occasionally misidentified as workers. Upon further analysis, a likely cause for this discrepancy is pinpointed as the domain shift between the scenarios of autonomous driving and our construction context. Contrary to the mid-range or long-range views typical in autonomous driving—where individuals are usually detected at distances exceeding 3.5m—workers in our setting frequently approach closer. Additionally, construction sites, unlike the open outdoor spaces prevalent in autonomous driving, are defined by their cluttered indoor nature. The scan profiles of various construction items might mimic human figures, and obstructions by these objects often lead to in-

[9]Here, "position and velocity" refer to the spatial and movement data within the map from a Bird-Eye-View (BEV) perspective. This clarification is important as, in much related literature, the position and velocity of a person are described in terms of their coordinates and movement on a 2D monitoring display.

complete scans of workers' profiles, further complicating the identification. These core domain variances likely induce a significant data distribution shift from the training datasets to the deployment environment. Addressing this challenge effectively requires a custom dataset specifically suited to our context and the realignment of models initially trained for autonomous driving applications. However, crafting such a dataset within a constrained timeframe exceeds the available capabilities of the research group[10]. Gathering a comprehensive set of LiDAR data depicting individuals among construction materials and laboriously annotating them in each frame to form a usable dataset demands an effort well beyond our available resources.

To circumvent this difficulty, the second iteration concentrates on exploring vision-based methods. Utilizing RGB-D cameras, the process involves first segmenting workers within the RGB image using bounding boxes and then estimating the 3D coordinates of the bounding box's center. If the center point of the bounding box is denoted as P_c, then its 2D coordinate on the RGB image (x_i, y_i) with respect to the center of the image can be easily determined through transformation. This transformation uses the coordinate of P_c with respect to the upper left corner of the image[11], (u, v), and the RGB image frame's half-length and width, x_c, y_c (Equation 4.1). By leveraging the depth information associated with the RGB pixels, the depth estimation of P_c in 3D space, z_p, is readily accessible. The task then simplifies to estimating x_p, y_p using basic projection relationships. With the camera's focal lengths denoted as (f_x, f_y), these can be calculated using Equation 4.2.

[10]In another study conducted by Haoyu Liu within the ReAC Research Group [Liu, 2023], similar challenges and findings were observed. Due to the complexity and unpredictability of real-world work environments, even well-tested computer vision algorithms often underperform and require custom fine-tuning. A significant focus of Liu's research was the use of generative and collective methods to create scenarios and datasets specific to real-world construction materials and objects, ensuring the algorithms are more aligned with actual conditions.

[11]This is the default origin setting of the selected RGB-D camera.

$$\begin{cases} x_i &= u - x_c \\ y_i &= v - y_c \end{cases} \tag{4.1}$$

$$\begin{cases} x_p &= z_p \frac{x_i}{f_x} \\ y_p &= z_p \frac{y_i}{f_y} \end{cases} \tag{4.2}$$

By default, the camera's intrinsic parameters, f_x, f_y, c_x, c_y, are known. Coupled with the camera's transformation relative to the robot's body, the workers' relative positions can be eventually determined. The crux of this pipeline, therefore, lies in selecting an effective vision-based object detection algorithm. Similar to the assessment of LiDAR-based methods, various versions of YOLO [Redmon et al., 2016, Wang et al., 2023] algorithms are tested[12] in the lab, with `YOLOv7` ultimately being selected. Unlike LiDAR-based methods, these algorithms exhibit relatively robust performance in recognizing people among diverse construction objects, mostly unaffected by the domain transfer issue previously mentioned. Though the FoV of an RGB-D camera is inevitably narrower than that of 3D LiDAR, as noted in the hardware design section, this approach greatly bypasses the cumbersome process of model retraining and alignment, offering a valuable trade-off.

Following the selection of this approach, several engineering measures are implemented to enhance the method's effectiveness in real-world scenarios. To counter the limited FoV and the decreased data quality at the RGB frame's edges, two RGB-D cameras are deployed to extend coverage across the robot's frontal area. Considering the paramount importance of

[12]These versions include `YOLOv2`, `YOLOv3`, `YOLOv7`, and their respective tiny versions, made available by `https://pjreddie.com/darknet/yolo/`

the robot's immediate front, an intentional overlap in the FoV is designed. This arrangement ensures that workers on the right front are detectable by both cameras, allowing the final estimate to synthesize outcomes from both, thus diminishing the impact of camera edge distortions. The technique used for merging the raw detections from both cameras, particularly for amalgamating duplicate detections, is DBSCAN [Ester et al., 1996], a widely used density-based clustering algorithm. The principle is that when two detections are within a certain proximity (i.e., below a set distance threshold, `eps` in DBSCAN), they are treated as a single detection, with the combined position estimate being the midpoint between them. This strategy relies on the assumption that the detection results from both cameras can arrive promptly within the same processing time frame, a requirement that introduces a challenge for computational efficiency. In practice, each camera's feed is analyzed by an individual `YOLOv7` model[13]. Consequently, two Convolutional Neural Network (CNN) models run concurrently on the robot during deployment. Essential engineering efforts are made to optimize and accelerate model inference via GPU. By addressing the computational bottlenecks found in the original software package[14], such as unnecessary tensor duplication and inefficient matrix manipulation, the final GPU utilization related to these two models is reduced to approximately 17%, a marked improvement from the pre-optimization performance (74%).

In addition to employing a dual-camera setup, empirical observations during laboratory experiments reveal that detection robustness can be improved when a person's face is within the camera's FoV. This discovery

[13]Previously, we experimented with using a single model to sequentially process inputs from both cameras. Due to model inference latency differences and the ROS communication mechanism, it is difficult to maintain a first-in-first-out (FIFO) sequence. For instance, detection results from the left camera would falsely be assigned to the right camera from time to time. Despite efforts to rectify this, the issue persisted, prompting the switch to a dual-model system, which proved to be a more straightforward and clean solution.

[14]`https://github.com/leggedrobotics/darknet_ros`

aligns with the understanding that the face constitutes a key visual identifier for a person. To leverage this insight and enhance detection capabilities, the cameras are deliberately angled upwards by 15 degrees. The adjustment ensures a more favorable vantage point for capturing workers' faces rather than legs. This straightforward yet impactful adjustment also contributes to improved detection when workers approach closer to the robot, compared to cameras positioned horizontally.

Besides these physical adjustments, two empirical thresholds are found beneficial for refining detection quality. Extensive physical experiments indicate that detections beyond 6m tend to be unreliable, as the depth data from RGB-D cameras is inherently noisier than that from 3D LiDAR. Therefore, in practice, detections farther than 6m are disregarded. Reflecting the "short-range navigation granularities" concern mentioned in Chapter 3, this decision prioritizes short-range detection accuracy and stability over broad long-range coverage. Additionally, it is observed that individuals within approximately 0.5m of the robot exhibit poor detection stability due to the low performance of the detection algorithm in cases of excessive occlusion, sometimes resulting in a loss of detection. However, at such close proximity, the integrated 2D LiDAR and 3D LiDAR already adeptly pick up the person in the laser scan. In such instances, the robot takes the person as an obstacle, coming to an immediate stop. After careful observation and analysis, this strategy of "treating a suddenly appearing person as an obstacle" is deemed suitable. Naturally, encountering someone unexpectedly close—whether for a human or a robot—elicits an immediate halt rather than a calculation of their possible trajectory. Hence, the 0.5m distance acts as a hidden safety threshold, establishing a nuanced safety protocol for close encounters with workers. For individuals detected beyond this threshold, their movements are assessed to inform socially aware navigation. Conversely, for those within this limit (whether due to sudden appearance or the robot's inadvertently

close approach), they are treated as obstacles, leading the robot to promptly stop to avoid collisions and ensure safety.

4.3.2 Tracking

The raw detection outcomes, despite employing threshold filtering and clustering, remain somewhat noisy for effective velocity estimation or subsequent RL inference. This situation can result in erratic or unstable estimation of workers' trajectories. To address this issue, Multi-object Tracking (MOT) methods [Luo et al., 2021] are crucial for refining detection results. These methods often utilize techniques such as Kalman Filters at their core, applying statistical approaches to estimate smoother, less noisy trajectories (concurrently including both position and velocity) of individuals' movements. In this stage, two algorithms are qualitatively assessed in the lab: a classical Kalman Filter-based SORT method [Welch et al., 1995, Bewley et al., 2016] and DeepSORT [Wojke et al., 2017]. Both algorithms deliver satisfactory outcomes in testing, with DeepSORT outperforming the classical method in scenarios involving severe construction object obstructions. However, this improvement comes at the cost of increased GPU usage and latency due to the computational demands of executing an additional DNN model inference. Given the sufficient reliability of the classical method, it is selected as the default for final deployment, with DeepSORT remaining as a readily available alternative. It is also empirically observed that the effectiveness of both methods is contingent upon the quality of the raw detection results. Inadequate detection inputs cannot be effectively compensated by either method to yield high-quality tracking outcomes.

The entire process of worker detection and tracking during real-time deployment is visualized and showcased in Figure 4.10. Outcomes at various stages are depicted using geometries of different colors. Raw position esti-

mations from the left (green) and right (red) cameras are illustrated with spheres. When the same individual is detected by both cameras, both a green and a red sphere appear in close proximity to the person. The synthesized estimation post-clustering is represented with blue cylinders. Finally, the filtered position (cyan cylinder) and velocity (magenta arrow) are displayed within the raw point cloud as the ultimate outcome.

4.4 Hierarchical Motion Planning

The motion planning framework for the robot adopts a hierarchical strategy that integrates reinforcement learning with search-based planning methods to effectively navigate through the cluttered site layout and busy carpentry workers. As highlighted in Chapter 2, recent RL-based social navigation methods show promising achievements for enabling fluid navigation in dense crowds, but they carry a hidden limitation due to the assumption of open-space layouts. This limitation is evident in two main aspects: (1) the inputs for these RL algorithms predominantly focus on the motion of nearby pedestrians, and (2) the demonstrations are largely conducted in spacious, empty rooms. Acknowledging the field's achievements and potential constraints, the applicability of these methods to our context-where static site objects consistently present significant obstructions-calls for further adaptation. Ideally, a lightweight approach that leverages existing efforts in open-space crowd scenarios, without starting from scratch and effectively addressing the unique challenges of our setting, is preferred.

The term "hierarchical" essentially implies a dual-layered approach (Figure 4.11). The first layer prioritizes safety and collision avoidance with both site objects and workers, leveraging established and effective research in classical collision avoidance methods rather than incorporating another RL

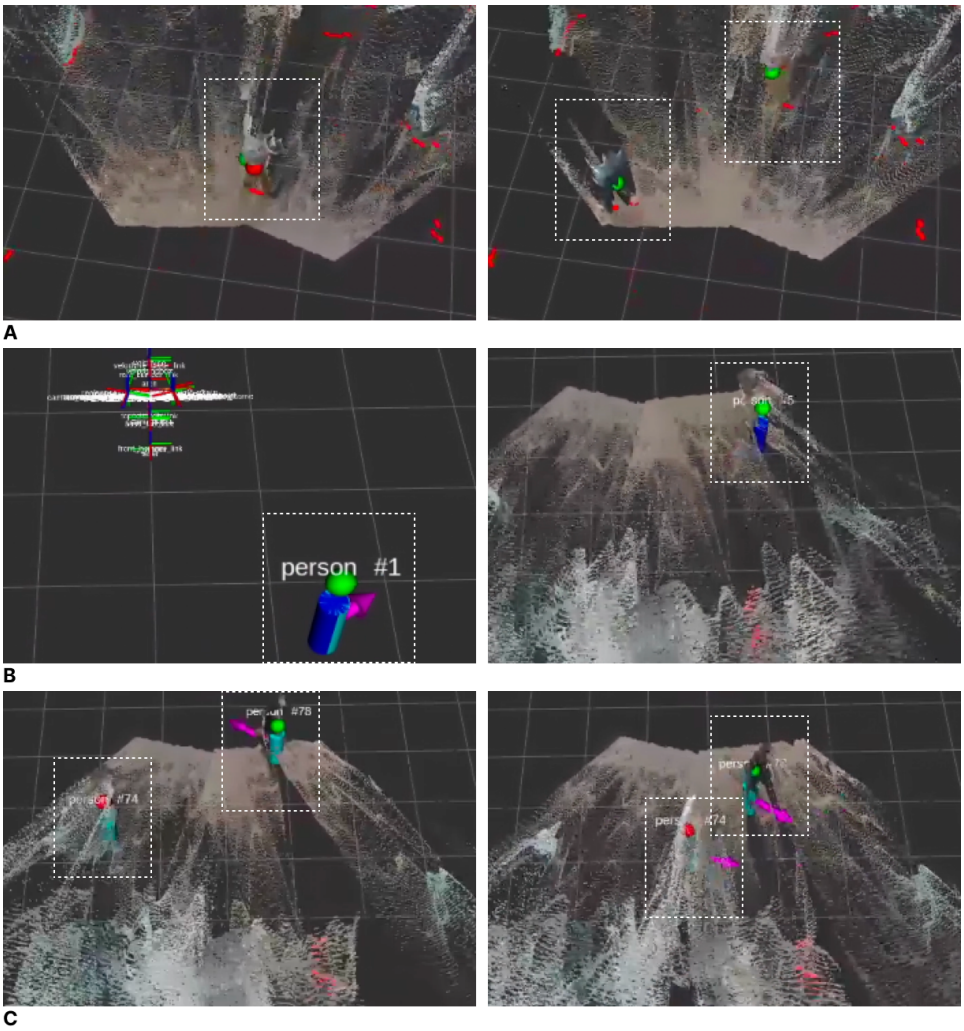

Figure 4.10: The entire worker detection and tracking process including (A) raw camera detections, (B) detection after clustering, and (C) the ultimate tracking result.

model. The second layer, focusing on social compliance and navigation comfort, employs RL-based social robot navigation algorithms. This divide-and-conquer strategy is informed by several practical considerations. Firstly, while overloading an RL model to handle both tasks is theoretically possible, it is neither strictly necessary nor efficient, given the plethora of proven methods available. Secondly, the layered structure is more computationally efficient through conditional activation; the RL layer is engaged only when workers are in proximity, reducing operation to the first layer in their absence. Since searching-based methods do not require constant DNN model inference, this results in low-latency and energy-efficient deployment. Third, the hierarchical design offers a generalizable and adaptable pipeline, allowing for experimentation with various classical collision avoidance algorithms and RL-based social navigation techniques with minimal engineering effort [Wu et al., 2024b]. Lastly, this strategy facilitates comparative analysis between classical and RL-based methods in real-world settings, offering greater explainability in assessing navigation performance and identifying potential issues through layer activation and deactivation, compared to relying solely on a "black-box" model.

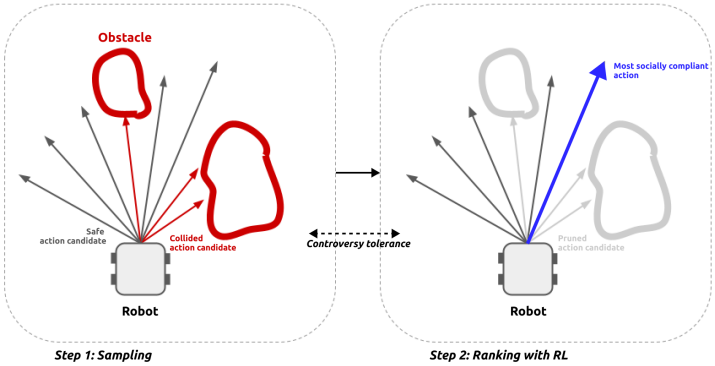

Figure 4.11: The dual-layered hierarchical motion planning stack.

Particularly, the continuous planar action space of the robot is first dis-

cretized into a pool of evenly spaced actions[15], denoted as $\mathcal{A}$. The aim of the first layer is to refine this pool by filtering out actions that are (1) not achievable due to the robot's velocity and acceleration constraints, or (2) likely to cause any form of collision in the next time step based on laser scan input. As a typical searching mechanism, although one could implement such filters from scratch, a more elegant solution involves leveraging what is readily available. These two filtering objectives align nicely with the velocity searching phase of the Dynamic Window Approach (DWA) [Fox et al., 1997], a classical method in robot navigation. The outcome of this phase, as shown in Equation 4.3, results in a refined action space, denoted as $\mathcal{A}_f$.

$$\mathcal{A}_f = \mathcal{A} \cap \mathcal{A}_{\text{admissible}} \cap \mathcal{A}_{\text{dynamic}} \tag{4.3}$$

Besides velocity searching, DWA also includes a trajectory evaluation phase, comparing trajectories unrolled by actions in $\mathcal{A}_f$ using a cost function. This function mainly considers three key factors: goal reachability, obstacle avoidance, and velocity maximization, essentially evaluating actions based on goal-reaching efficiency. This criterion suffices in the absence of workers. However, when movements of workers are detected by the upstream worker detection and tracking module, the emphasis shifts from goal-reaching efficiency to social awareness and comfort as the more critical criteria for choosing the robot's actions and trajectories. Thus, when workers are present, the hierarchical motion planning framework bypasses DWA's evaluation phase, allowing RL-based social navigation algorithms to take over the evaluation process.

To present a feasible cost function for evaluating an action's social com-

[15]In related literature, these actions are also called motion primitives

pliance, the second layer leverages value-based or actor-critic typed social navigation methods. Both approaches involve training a state value network[16] $V(s)$. In reinforcement learning, a state value function by definition estimates the expected cumulative rewards of a given state. In simpler terms, when trained and optimized, it aids in identifying the optimal policy and can serve as a "scoring" criteria to determine the optimal action in the next step. Based on the Bellman Equation [Sutton and Barto, 2018], such a process can be captured in Equation 4.4[17]. In social navigation, the reward function (see details in the next section) is tailored towards social awareness around people while approaching the goal. Therefore the value network based on this reward can be used to compare actions with respect to their social compliance and comfort among workers, complementing the aforementioned evaluation process.

$$\mathbf{a}_t^* \leftarrow \text{argmax}_{a_t \in \mathcal{A}_f} R\left(\mathbf{s}_t^{\text{joint}}, \mathbf{a}_t\right) + \gamma^{v_{\text{pref}} \cdot \Delta t} V_{t+\Delta t}^{\text{joint}} \tag{4.4}$$

In a practical sense, the engineering essence of this hierarchical motion planning stack is the "handover" of the refined action space $\mathcal{A}_f$ from the first layer to the second. Given the computational complexity differences, without specialized management, the action searching pool assumed in DWA, even post-refinement, tends to be larger than that typically assumed in RL-based social navigation research. Therefore, prioritizing computational efficiency and minimizing potential latency is crucial for their collaborative operation. The number of actions in $\mathcal{A}_f$ during a relatively smooth DWA-

[16] In the current trend of deep reinforcement learning, the value function is often approximated by a deep neural network. Therefore, it is referred to as a value network rather than a value function here.

[17] One may notice that this equation seems to be missing the model term $P(s \mid s, a)$. Following a convention in related literature [Chen et al., 2017, 2019], it is assumed that workers will continue their current velocity in the immediate next short time interval Δt.

based navigation can range between 80 to 120. Conducting value network inference for all these actions consistently within each decision timestep[18] necessitates essential computational optimization. By aggregating these individual actions into a 2D tensor and utilizing GPU capabilities, batch inference accelerates the process.

4.5 Contextualized Fine-tuning of RL

With the hierarchical motion planning stack serving as a foundational structure, the application of reinforcement learning (RL) focuses on enhancing the robot's ability to adapt and navigate around workers in a socially considerate manner during their active work duties [Wu et al., 2024b]. This component of the research integrates all preceding modules and serves as a critical junction where the previously identified challenges, nuances, and specificities of construction work are collectively addressed and resolved. Drawing upon crucial insights from earlier chapters and sections, the key challenges for RL-based social navigation can be summarized as follows: (1) while the worker count within the target environment generally does not exceed five, the complexity, diversity, and intensity of their activities are significant; (2) as elaborated in Chapter 3, the behavior patterns of workers deviate substantially from those of pedestrians, which often serve as the foundational assumption in much of social navigation research; (3) considering the established hardware configurations in the Hardware System Design section, worker detection primarily occurs within the robot's immediate frontal area, challenging the often-simulated presumption of the robot's omniscient, mostly perfect, 360-degree awareness of surrounding agents in common symbolic simulation environments, thereby necessitating specific adaptations to this reality.

[18] Practically, the system's decision timestep is set to 0.1s on ROS.

With these high-level observations as the foundation, a pivotal decision before proceeding is determining the research focus—specifically, whether addressing context-specific challenges involves reinventing an algorithm or adapting existing methods. This decision is informed by two key insights. The first insight draws from recent advancements in RL-based social navigation research [Liu et al., 2021, Wang et al., 2022], where the efficacy of popular algorithms has been benchmarked, exhibiting relatively good navigation performance in scenarios with around five agents. These assessments are conducted in both simulated and physical environments using simplistic setup patterns like square and circle crossings. Such configurations, where agents[19] navigate from one side of a geometric shape to crossing points on another, effectively mimic dense pedestrian dynamics. However, the scenarios in which social navigation capabilities are demonstrated, including both the modeling of agent behaviors and the setup of the environment, cannot fully capture the diversity and complexity of worker behaviors and activities in a construction setting. How the inclusion of these factors might affect the performance of these algorithms remains an open question.

The second insight pertains to the common use of ORCA [van den Berg et al., 2011] during the early phases of RL model training as a means to expedite learning. ORCA's limitations——such as oversimplifying human behavior, neglecting social norms, and struggling with dense interactions——can potentially carry over to these trained RL models if further refinement and adaptation are not applied. Recognizing these existing efforts and their limitations in addressing the specific needs of our context, the objective of this module is not to develop an entirely new RL algorithm, but rather to design a generalizable fine-tuning and alignment pipeline that builds on the generic social navigation capabilities of existing methods and improves them to ef-

[19]In social robot navigation literature, "agent" is commonly used to refer to either robot or human. Here, "agent" specifically denotes symbolic entities within the simulation environment.

fectively address the unique challenges and context-specific nuances outlined above.

Following this decision, the research proceeds by selecting a performant, existing RL-based social navigation method as a base for future alignment, guided by both literature review and qualitative laboratory testing. The candidate algorithms evaluated include SA-CADRL [Chen et al., 2017], GA3C-CADRL [Everett et al., 2018], LSTM-CADRL [Chen et al., 2019], and SARL [Chen et al., 2019]. As mentioned in the previous section, an important criterion in drafting this pool of algorithms is that they are either value-based or actor-critic typed methods, each involving the training of a value network. Also adhering to a similar Markov Decision Process (MDP) framework, these algorithms are compatible with our framework, facilitating straightforward deployment, experimentation, and comparative study on the robot through the simple exchange of the inference model file. While in theory, any of these methods could serve as a foundation for subsequent RL model adaptation, SARL is ultimately chosen due to its relatively stable performance in preliminary robot testing. Qualitatively, when interacting with humans, the SARL model demonstrates a notably smoother response, strategically planning circumventing trajectories well in advance, as opposed to executing abrupt turns at close distance or reactions too late.

Based on this practical selection, the following subsections primarily focus on how a pre-trained SARL model is contextually fine-tuned and aligned by incorporating multiple key factors. These factors, such as worker behaviors, layout setups, and limited field of view, critically articulate the domain differences between generic crowd scenarios and our context and are fused into the adaptation of the RL model through customization in the simulation environment. To form a uniform adaptation scheme among these diverse factors, the section concludes with a generalizable, efficient curricu-

lum learning pipeline that can enable incremental fine-tuning and testing of the model, especially concerning its continuous deployment on the robot.

4.5.1 Customized Agent Behavioural Patterns

To bridge the gap between pedestrian behaviors typically assumed in social robot navigation research and the diverse, work-oriented behaviors of carpentry workers, four specific agent patterns are devised in the simulation environment. In simple terms, an agent pattern refers to a predefined way a robot or virtual worker is expected to move or act in a given situation. These generative patterns, alongside the conventional collision-avoidance walking pattern used in most related studies, provide a more nuanced reflection of the uncertainty and variability encountered in a construction setting. Insights from Chapter 3, particularly the analysis of individual work paradigms (Figure 4.12), such as working adjacent to the formwork wall, walking/traversing, and standing still, inform the design of these patterns.

- **Pattern I, Stop-and-Go**: This pattern emulates the intermittent work rhythm of carpenters working by the wall, characterized by periods of intense activity followed by brief movement to the next spot. The agent's velocity intermittently drops to zero, simulating a pause, based on a time interval sampled from a Gaussian distribution. Concurrently, the agent's radius varies significantly to mimic the worker's expansion of personal space during work phases.

- **Pattern II, Walking While Carrying Things**: Reflecting the traversing or walking paradigm, this pattern accounts for the unique challenges workers face while transporting materials and tools. Despite a superficial resemblance to pedestrian motion, the safety radius of these agents is adjusted to accommodate the extra space required

for carried items. This pattern also incorporates a Gaussian process to model moderate fluctuations in the safety radius. Notably, agents in this pattern do not exhibit collision avoidance behavior, simulating the restricted visibility and maneuverability workers experience when carrying panels or tools.

- **Pattern III, In-Place Fluctuation**: Designed to represent workers who are stationary but engaged in tasks requiring movement within a confined space, this pattern deviates from the notion of complete stillness. Instead, agents exhibit slight positional shifts within a designated circular area, simulating the natural movement of workers attending to installation tasks in a relatively fixed location. This behavior, similarly along with a fluctuating safety radius modeled through a Gaussian process, reflects the dynamic nature of stationary work tasks and the spatial requirements of the worker's tools and materials.

- **Pattern IV, Back-and-Forth**: This pattern simulates workers moving back and forth, intermittently switching goal positions between two points. It reflects the common movement of Worker B during panel preparation, and Worker A's movement along the target wall for miscellaneous checks and installations. Apart from the shifting goal positions, the agent's safety radius and velocity are mostly similar to those of a pedestrian, reflecting the relatively straightforward nature of this movement pattern.

4.5.2 Customized Agent Layout Configurations

Reflecting the collaborative work paradigms and the confined, dense nature of construction work activities, the simulation environment's experiment configurations are customized (Figure 4.13). The simulation space is de-

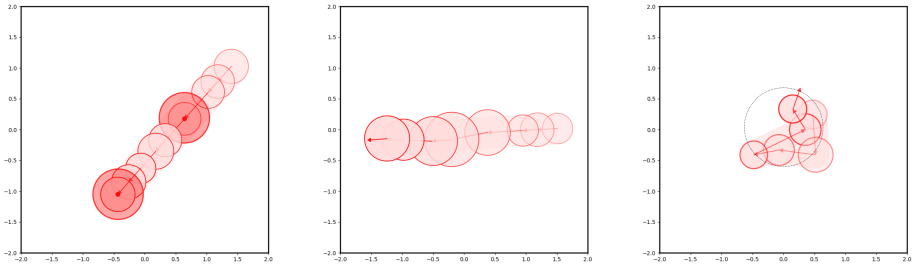

Figure 4.12: The customized agent patterns.

liberately limited (2.5-3m) to replicate the constrained operational areas in carpentry formwork, requiring agents to navigate within much tighter confines. Experiment layouts are devised by creatively employing the customized worker behavioral patterns and incorporating insights from the two collaborative work paradigms. Additionally, scenarios specifically designed to provoke close encounters with the robot are introduced. These scenarios strategically place agents in proximate positions to the robot from the onset, diverging from the more expansive, distant crossing scenarios commonly employed. This design aims to evaluate the model's ability to execute immediate, sharp maneuvers in abrupt encounters—a capability not extensively explored in existing research. Each test case is established through a rule-based, generative approach, allowing for dynamic variation rather than confinement to a predefined set of positions. For ease of reference, the following description directly mentions the customized agent behaviors using Patterns I, II, III, and the collaborative work paradigms analyzed in Chapter 3 as Paradigms I and II.

- **Case I, Work Interaction**: Formulated around the interaction-focused collaborative Paradigm I, this scenario generates two agents. The first agent, inspired by Worker A, displays behaviors from Patterns I and III. The second agent, simulating Worker B, can adopt

any of the three patterns at random. Initial positions for these agents and the robot are randomized within the compact area, avoiding rigid geometric layouts to better mimic the unpredictable nature of human-robot interactions in constrained spaces.

- **Case II, Work Separation**: Drawing from the separated work activities of collaborative Paradigm II, this scenario introduces three to five agents. One agent represents Worker A and may randomly exhibit any of the three behavioral patterns. The rest are equally likely to follow Pattern II—reflecting workers in transit with tools and materials—or act as typical pedestrians with collision avoidance capabilities, resembling Worker B or others momentarily traversing the site. This scenario offers a more open spatial arrangement (6m) than Case I.

- **Case III, Front Encounter**: This setup involves 2 to 3 agents, with one positioned to cross directly between the robot and its goal at the start. Similar to Case II, the other agents either follow Pattern II or act as generic pedestrians. To navigate past this initially close agent, the robot must immediately adapt its path, showcasing its agility and strategic decision-making.

- **Case IV, Side Encounter**: Mirroring Case III's dynamics but with a lateral approach, this scenario positions 2 to 3 agents to pass closely by the robot's side initially. To enable socially adept interaction, the robot is required to quickly adjust its trajectory to keep a respectful distance, testing its prompt responsiveness and the sophistication of its navigation algorithms during close encounters.

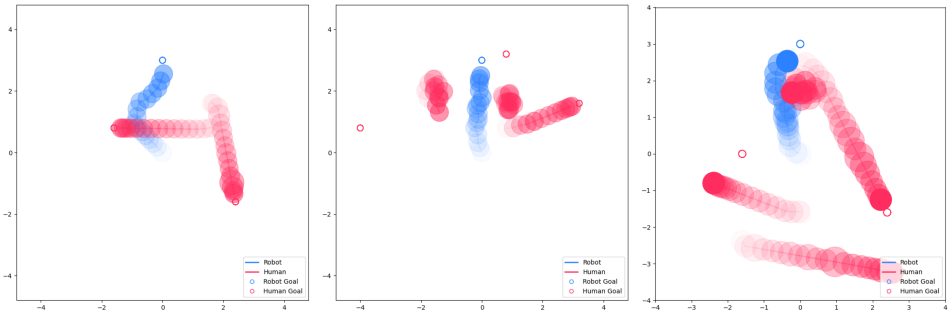

Figure 4.13: The customized agent layout configurations.

4.5.3 Limited FoV of the Robot

Rather than a full, 360-degree coverage of all agents, the robot's practical Field of View (FoV) is limited by the RGB-D camera setup, covering approximately 110 degrees of the robot's frontal space. To account for this partial observability, the simulation environment incorporates a blind zone around the robot, intentionally omitting agents that appear within this zone from the robot's observation. This approach, drawing insights from related literature [Stoler et al., 2022], introduces a level of partial observability, sometimes referred to as First-Person-View (FPV). Although this can impact the RL model's performance by reducing its field of awareness, it more accurately reflects real-world conditions and challenges, aligning the simulation environment closer to the practical scenarios the robot will encounter on site.

4.5.4 Domain Randomization

To enhance the Sim2Real transition and adapt the model to imperfect inputs derived from raw sensory data, a domain randomization scheme has been developed. This includes temporarily introducing occlusions to the simulated agents' states (position, velocity, safety radius), simulating instances of temporary detection loss akin to what might be encountered with

real-world sensory inputs. Additionally, random noise is deliberately added to the agents' states before they are fed into the model, purposefully crafting imperfect inputs for the model. These measures aim to make the RL model more robust to the kinds of input variability and imperfections it will encounter in real-world scenarios, thereby improving its performance and reliability when deployed on the robot.

4.5.5 A Uniform Pipeline For Incremental Model Adaptation

To integrate these diverse factors into the model through a cohesive approach, the decision is to employ curriculum learning. Viewing all factors as elements of a progressively deepening curriculum, the pre-trained vanilla SARL model, initially optimized in a generic setting, undergoes successive refinement through the series of customized experimental scenarios. This methodical approach is chosen not only for its structural uniformity but also due to some critical advantages it presents. Most notably, it facilitates the incremental enhancement of the same model, obviating the need to start from scratch for each new adjustment. This feature, while potentially overlooked when all training occurs within a simulation environment, proves exceedingly valuable for bridging the gap between simulation and real-world application (Sim2Real) and conducting physical tests on the robot. It significantly mitigates the risk of unpredictable or radically altered robot behaviors emerging during the rapid adaptation of the model. The aim is to preserve a consistent level of functionality for previously tested robot capabilities while systematically integrating new features and improvements. This approach eschews the pitfalls of testing entirely new models from the ground up for each adaptation, which can be both unstable and inefficient.

This necessity calls for measures to mitigate potential volatility in the

RL model's performance due to significant policy shifts during model adaptation. Specifically, when fine-tuning the value network via gradient descent, it is crucial to ensure that the model, post-training, does not substantially deviate from its pre-training state. In canonical policy-based methods, this stabilization is typically achieved by parameterizing the policy to follow a well-defined probability distribution, such as the Gaussian distribution. This allows for an easy estimation of the shift between the post-training policy, π_{new}, and the pre-training policy, π_{old}, using KL-divergence, a measure for the difference between two probability distributions. However, in our case, making such an assumption is baseless, especially considering that our chosen RL algorithms can be value-based.

To address this challenge, we draw inspiration from measures taken in RLHF [Christiano et al., 2017], estimating the KL-divergence of policy shift through an alternative method. Since the value network in reinforcement learning reflects the policy's main features, we indirectly estimate policy shifts through the changes between post-training and pre-training value networks. Given that the robot's action space has been discretized and refined in previous steps, this shift can be particularly gauged by comparing the two value networks' preferences for the same set of actions $\mathbf{a}_1, \cdots, \mathbf{a}_m$ within the refined action space $\mathcal{A}_f$. Using each action's inferred value as a score, this preference is measured utilizing the Plackett-Luce model [Plackett, 1975, Luce, 2005] (Equation 4.5), effectively transforming value rankings among actions into a probabilistic distribution (Equation 4.6), thus enabling the efficient estimation of KL-divergence (Equation 4.7).

To constrain model adaptation and penalize policy shift effectively, the derived KL-divergence is added to the original reward function r_{ori}, as introduced in [Chen et al., 2019], serving as a penalization term. Thus, while we encourage the model to adapt to training curricula, we also penalize it

for making overly drastic changes. This strategy ensures that the model's adaptation occurs gradually and stably, ultimately minimizing the risk of abrupt behavioral changes in the robot that could destabilize its real-world performance.

$$P(\mathbf{a}_i|V) = \frac{V(\mathbf{a}_i)}{\Sigma_{j=1}^{m} V(\mathbf{a}_j)} \tag{4.5}$$

$$D_{\mathrm{KL}}(\pi_{\mathrm{old}} \| \pi_{\mathrm{new}}) = \Sigma_{i=1}^{m} P_{\mathrm{old}}(\mathbf{a}_i) \log \left(\frac{P_{\mathrm{old}}(\mathbf{a}_i)}{P_{\mathrm{new}}(\mathbf{a}_i)} \right) \tag{4.6}$$

$$r = r_{\mathrm{ori}} - \lambda D_{\mathrm{KL}} \tag{4.7}$$

In summary, this structured methodology, encompassing (1) experiment setup customization, (2) a curriculum training pipeline, and (3) a theoretical basis for mitigating model adaptation instability, collectively forms a comprehensive strategy for addressing the unique challenges encountered in a real-world, populated and complex work scenario. Notably, this approach—characterized by tailored datasets and experiments, a sequential curriculum, and stable, continuous model integration—can also serve as a generalizable pipeline for fostering consistent, incremental performance enhancements in AI-driven robotics beyond the current context.

Concluding this chapter, Figure 4.14 presents a system dashboard that encapsulates the entire architecture of the robotic framework. This dashboard elucidates the engineering intricacies of how the various modules are interconnected and orchestrated within the live robot prototype, providing a visual articulation of the system's complexity and organization.

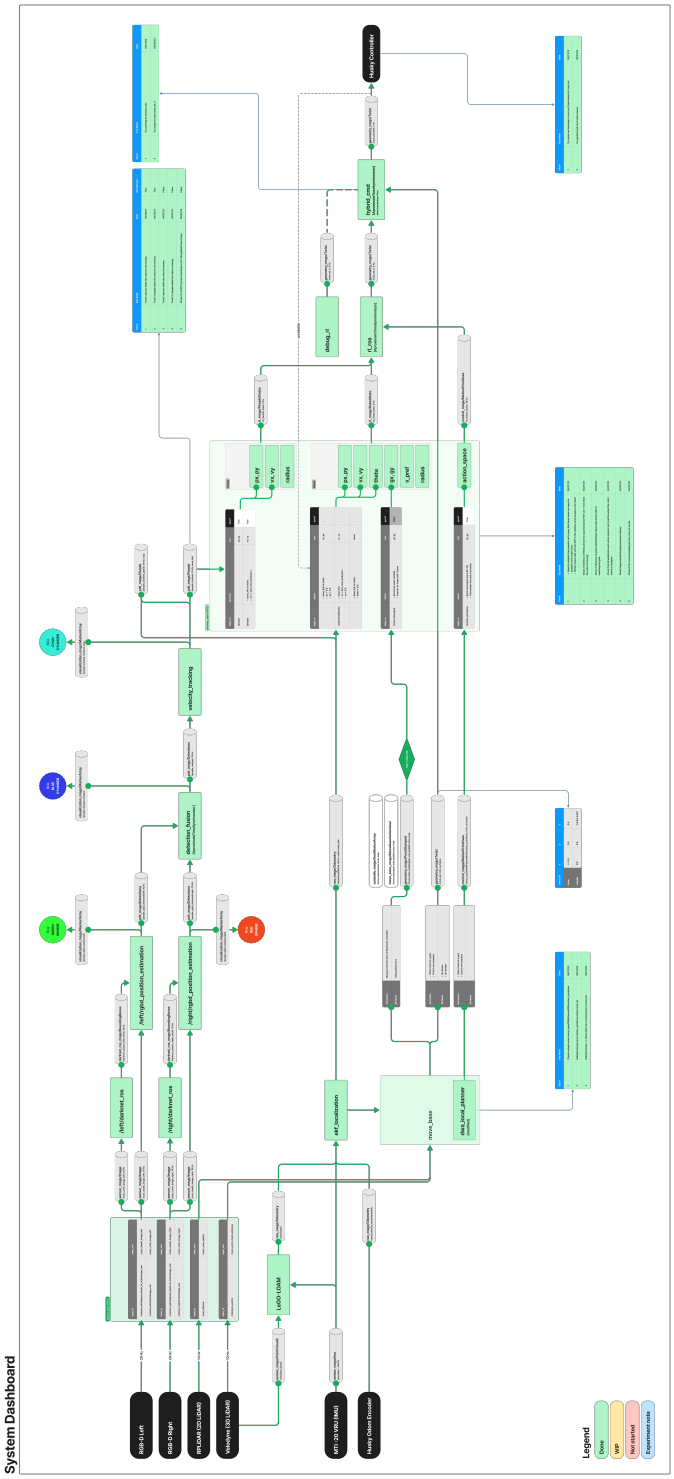

Figure 4.14: Overall system dashboard of the robotic framework. For further clarity, a larger version of this image can be accessed online: `https://ice-5.github.io/assets/img/thesis/dashboard.png`.

Chapter 5

Results

This chapter presents the key results of this research, beginning with an overview of the realized robot prototype. The main outcome demonstrates the effectiveness of the proposed framework in enabling the robot to navigate and assist in dynamic, real-world construction settings. Qualitative and quantitative evaluations were conducted to assess the robot's performance in two primary scenarios: (1) an active construction site with workers, and (2) a lab-based sandbox environment reflecting site features. The results indicate that the robot can provide socially aware and comfortable support to workers, adapting its behavior to site-specific challenges.

During the on-site evaluation, the robot prototype was deployed on an active construction site, where its performance was qualitatively examined through direct interactions with workers and follow-up interviews[1]. Workers' feedback and on-site documentation indicate that the envisioned robotic support can function within a real construction site, with the RL-driven social navigation capabilities proving responsive among workers. The evaluation also identifies areas for improvement and further development in the

[1]The evaluation was carried out in a designated area under close supervision by research group members during worker interactions.

current framework.

In the lab-based sandbox evaluation, the robot's support capabilities were quantitatively tested across a series of navigation and interaction metrics, within layouts that reflected key environmental challenges observed during the on-site demonstration. Two comfort-related metrics—comfortable delivery rate and comfortable encounter rate—were developed to evaluate the robot's ability to provide tangible and considerate assistance to workers in close proximity. The quantitative evaluation validated the effectiveness of reinforcement learning in enabling socially aware and comfortable support. Additionally, it showed that the contextual fine-tuning of the RL model led to smoother navigation performance.

5.1 Final Robot Prototype

Figure 5.1 illustrates the final state of the robot prototype, showcasing both the key sensor and hardware components, along with the adjustments made for on-site deployment. Beyond the core technical features detailed in Chapter 4, the build incorporates structural fortification, wiring organization, battery and heat monitoring, and port management. While these efforts might seem minor, they are crucial for maintaining the robot's operational integrity. The construction site floors, often cluttered with large debris and gravel, can significantly jostle the robot's upper structure, particularly affecting the 3D LiDAR box, RGB-D camera stands, and the numerous USB ports connecting various sensors. These connections frequently loosen without notice, leading to a range of inexplicable issues. To mitigate such risks systematically, every connection is secured with hot glue, including both mechanical fastenings and USB connectors. Additionally, the robot's battery

occasionally overheats[2], necessitating continuous monitoring both from the robot's chassis and its onboard computer. Lastly, to ensure safety around construction workers, all sharp metallic corners, edges, and protrusions are covered with bubble wraps and brightly colored duct tape.

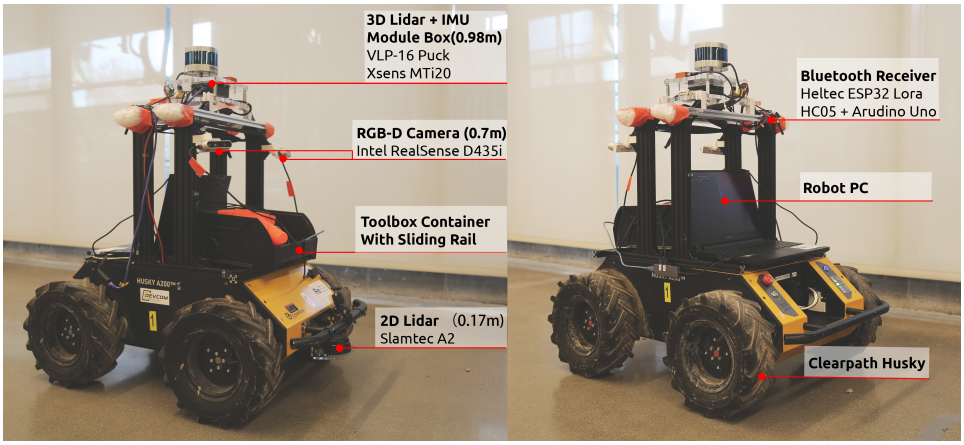

Figure 5.1: Final robot prototype. This image was previously published in [Wu et al., 2024b]

5.2 Simulation Training Results

Figure 5.2 shows the pre-training and fine-tuning of the RL model in the simulation environment, focusing on cumulative reward and collision rate. While assessing the robot's navigation comfort around workers, along with workers' reactions and feedback, is challenging in a simulation environment, the training process demonstrates the model's adaptation as it is fine-tuned in context-specific scenarios. The reward initially drops when the agent's behavioral diversity, as described in Chapter 4, is introduced but gradually recovers and stabilizes at a consistent level, similar to the pattern observed in the collision rate. This indicates the RL model's ability to adapt to the new context, accommodating the introduced agent and scenario diversity

[2]One battery completely burnt out due to this issue.

and uncertainty. The fine-tuned RL model is then deployed to the robot prototype for further human-centric evaluation, both in the lab and on site.

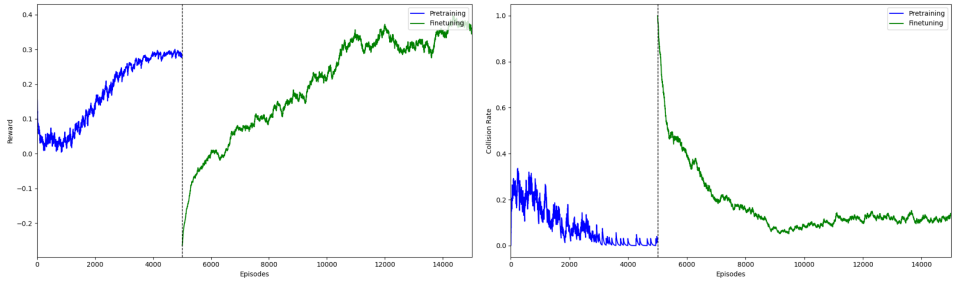

Figure 5.2: Training process of the contextual fine-tuning of the RL model in the simulation environment, shown by cumulative reward (left) and collision rate (right). A similar version of this image appeared in [Wu et al., 2024b].

5.3 Qualitative Evaluation On Site

5.3.1 Site Setup

The on-site demonstration and evaluation of the robot prototype take place on an active construction site on campus, within a multifloor facility building. During the initial field tour, three potential areas were presented to the research group as tentative candidates[3]. In terms of construction phases, two areas are in the interior phase, subsequent to carpentry formwork, with concrete structures completed and thin steel beams installed for non-load-bearing room partition walls. The third area is in a phase roughly preceding interior detailing, designated to be a spacious, column-free court space. This space does not contain dense room partitioning steel beam arrays and is

[3]The on-site demonstration would not have been possible without the extensive administrative efforts in coordinating with our industrial collaborator for sites, personnel, permissions, etc. Given the dynamic nature of construction sites and the day-to-day progression, it is infeasible to find a site where carpentry formwork installation is actively ongoing. Consequently, the goal is to identify a site with similar layout complexities and construction object distributions to those of a carpentry formwork site, among workers knowledgeable enough to simulate carpentry formwork installation workflows.

temporarily used for material storage, occupied by clusters of construction material piles. Among these three options, the research group deems the third area to more closely resemble the geometric and spatial features of a carpentry formwork site, since the dense array of thin steel beams in the other two areas would significantly alter the spatial layout and navigation challenge[4].

Consequently, a portion of the third court area, measuring roughly 45m x 35m, is designated as the demonstration and evaluation area (Figure 5.3). This setup includes a target formwork wall and several sequential work zones along it. Approximately 15m from the wall and work zones, a cluster of toolboxes and material piles is assigned as the central workbench[5], where tools and materials for carpentry formwork are stored and dispatched. Apart from these assignments, the demonstration area is mostly left in its original state, with construction workers continuing their daily activities. Both the work zones and central workbench are designated in a general manner, without specific coordinates or QR codes. Additionally, the site is not intentionally tidied or cleared to ease robot navigation, ensuring the evaluations remain close to real-world construction work conditions.

5.3.2 Worker Interaction Scripts

To better communicate with workers regarding the process of the robot demonstration, three demo scripts are formed to facilitate the briefing and explanation. These scripts capture key carpentry formwork activities, as well as the major robot support capabilities described in Chapter 3, particularly

[4]Once the steel beam arrays are installed, it would be preferable to consider a smaller robot due to the transition from an open layout to a roomed layout. Navigation in such a layout involves not only collision avoidance but also path planning among rooms, considering room connectivity as a prior. This presents an additional layer of challenge not foreseen or studied in the carpentry formwork context.

[5]This assignment takes into account the distance estimation and layout features observed in previous site studies of carpentry formwork.

Figure 5.3: On-site demonstration and evaluation setup. The photo was also covered in [Wu et al., 2024b].

testing the robotic functions of **sending**, **summoning**, and **accompanying**. As illustrated in Figure 5.4, the scripts are as follows[6]:

- **Script 1: Send**: The robot's initial position is set at one of the work zones along the target wall. Workers send the robot to the central workbench to retrieve tools and materials, while they continue their work on the target wall. The sending can be initiated by any one of the workers by pressing button `1` (representing the central workbench) and `confirm` on the robot's keypad controller.
- **Script 2: Summon**: This script follows the previous script. The robot's initial position is at the central workbench, where it is loaded with tools and materials. Workers summon the robot to return to their work zone along the target wall by pressing button `2` (encoding the current work zone) and `confirm` on the robot's controller. Notably, this summoning can be initiated from a distance, not necessarily close to the robot.
- **Script 3: Accompany**: Workers wrap up their work at the current work zone and move to the next work zone along the target wall. The robot's initial position is in the current work zone. The workers place heavy tools and materials into the robot's container and call the robot to accompany them to the next work zone by pressing `3` and `confirm` on the controller. The robot then follows the workers to the next work zone, maintaining an appropriate pace and social awareness.

An additional test case is designed to challenge the robot's long-distance

[6]The way this research documents and groups these scripts is for the purpose of technically analyzing the robot's performance, particularly echoing the robot's proposed support functions mentioned in the previous chapters. One may notice other ways of categorizing these scripts in ReAC's other reports and publications, for instance, some grouping Script 1 and Script 2 as a single script. However, the way of grouping does not fundamentally change the actual content of the analysis.

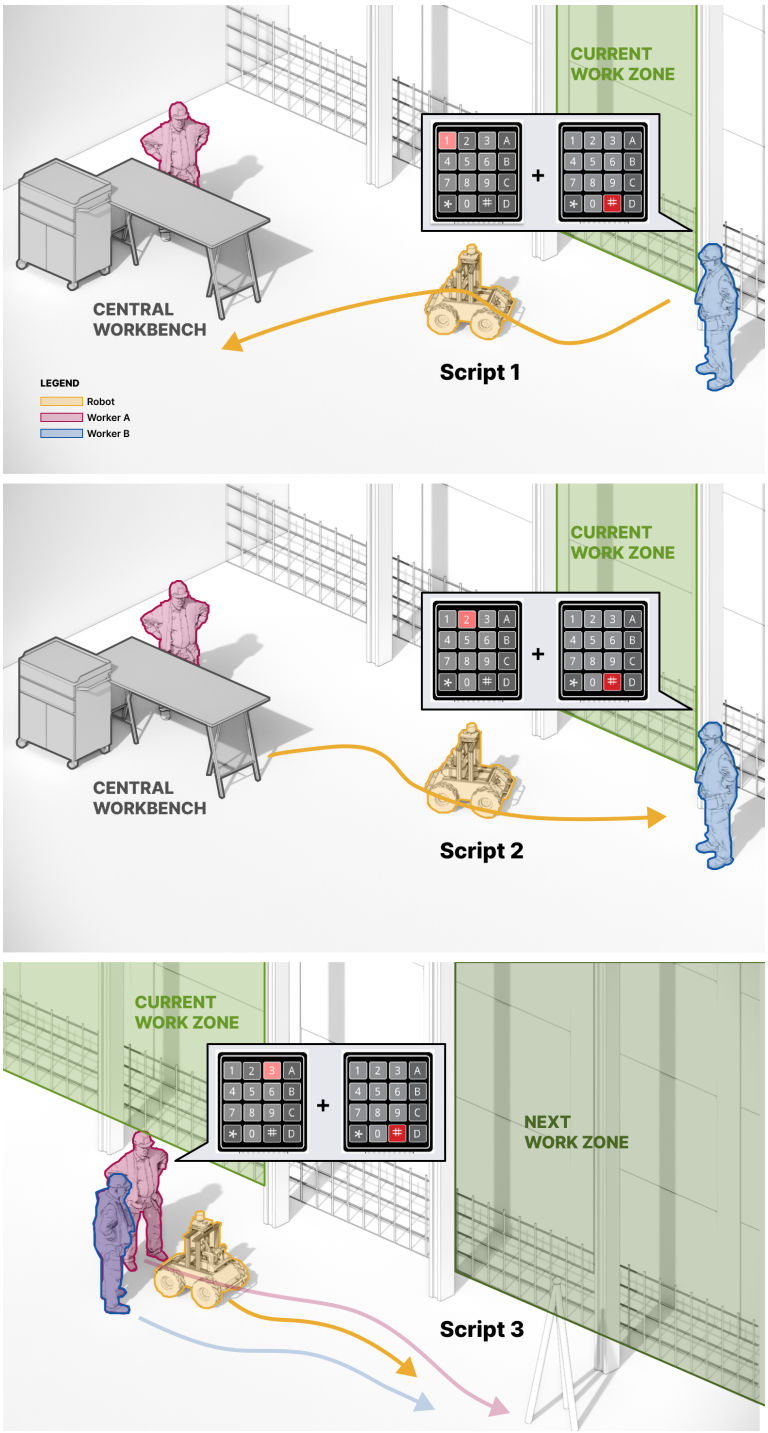

Figure 5.4: Storyboard of the three scripted scenarios.

navigation robustness and adaptability in an unfamiliar, sparsely mapped area. This challenge, while less relevant to the main support functions, is designed to uncover potential weaknesses and corner cases. It focused on identifying issues in the robot's navigation capabilities on site.

5.3.3 Pre-Demonstration Preparation

A day before the demonstration, the robot is transported to the site to conduct preparatory work. This includes the initial mapping of the demonstration area, as well as a series of preliminary tests to ensure the robot's operational readiness. In the mapping process, research group members drive the robot around the designated demonstration area using a joystick to create a preliminary map of the site (Figure 5.5). Although the material piles and arrangement of tools might change due to ongoing worker activities, this mapping aims to prepare the robot with a baseline layout of the site. Both 2D mapping (`gmapping`) and a customized mapping method based on LIO-SAM are conducted.

Using the base maps, group members also drive the robot around the site to encode key locations and areas. This encoding is done by driving the robot using a joystick to the intended place and pressing the buttons on the keypad controller to record. This convenience is underpinned by the robot's mapping and localization module, which memorizes the robot's state relative to the base map.

After these procedures, group members simulate the demonstration scenarios by acting as workers and conduct 10 runs per script. The robot's battery is fully charged, and the onboard computer is booted up to verify the ROS system's stability. The robot sensors are double-checked for robustness and structural stability on the rugged construction floor. The

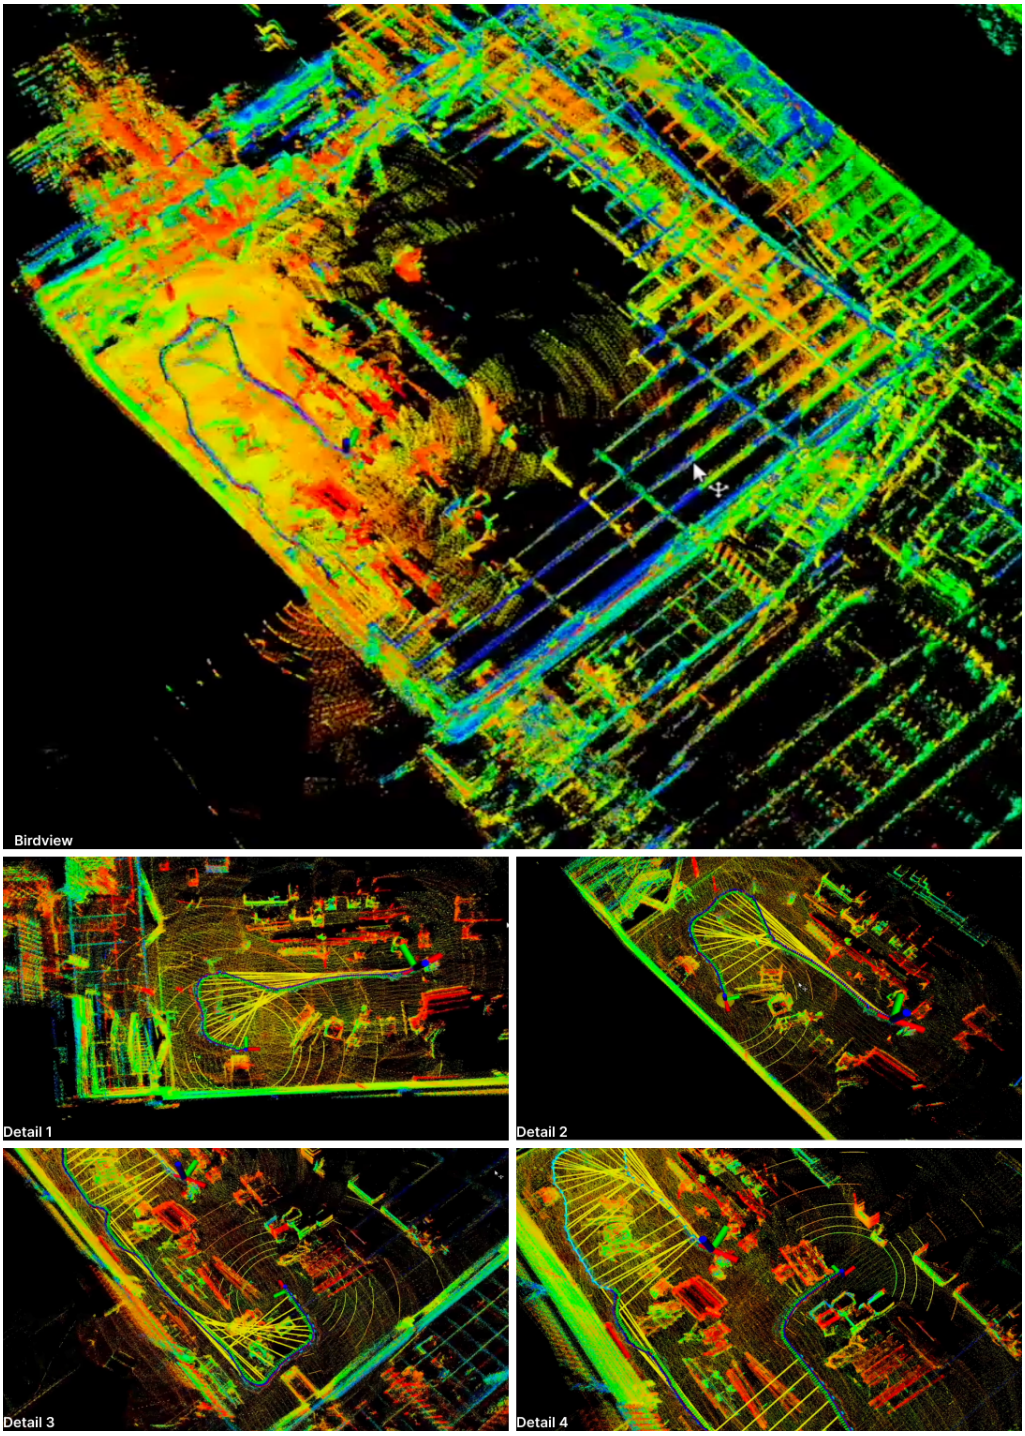

Figure 5.5: Mapping of the demonstration site. Part of this image was included in [Wu et al., 2024a].

navigation stack is tested for performance and integrity in the site's dynamic environment. The group also conducts a final check on the robot's hardware, including the 3D LiDAR, RGB-D camera, and the custom-built keypad controller. The robot's safety features are reviewed, including the emergency stop button and the safety bubble foam covering sharp edges. Additionally, four time-lapse cameras are set up around the site to capture the robot's movements and interactions with the workers during the demonstration.

During the final check, an issue related to horizontally placed, elevated, protruding pipes is noted (Figure 5.6, left). Although the LIO-SAM based mapping process can capture these pipes with the 3D LiDAR, they are not promptly detected in real time during navigation due to small blind areas uncovered by the sensors. Thin pipes or steel beams protruding at approximately 0.35 to 0.5m above the ground posed significant challenges. This failure to capture these thin elevated objects in real time can lead to instability in both localization and local path planning. In terms of localization, a failure in AMCL can cause drifts in state estimation if these structures are not captured. For local path planning, which relies heavily on the scan input for the local costmap, if these objects are not captured or have moved to a different location compared to what is recorded in the global costmap, then the robot would have no viable input about these risks and can easily collide with them. This corner case, not observed in the preliminary site study, is noted frequently during pre-demonstration testing.

To quickly address this issue, the intuitive fix is to amend and cover the small blind area among the current sensors. For this purpose, the research group installs an additional RGB-D camera in the front of the robot, attached on a stand placed above the robot's front bumper (Figure 5.6). Given the camera's noisy, ranged raw point cloud input when directly con-

verted as scan input, a simple distance-based filtering is applied to this source[7]. This additional camera is then rapidly integrated into the robot's navigation stack, with a few validation tests conducted overnight in the lab.

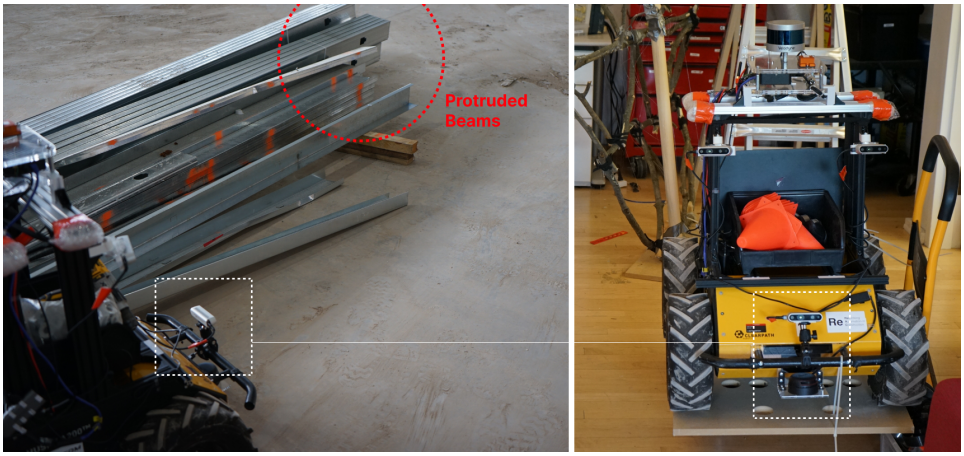

Figure 5.6: Rapid installation of the third camera in the lab.

5.3.4 Robot Demonstration Performance

Before and during the demonstration, research group members brief the participating workers on the details of each script and their roles in the evaluation. The robot is then deployed to the work sites, and the three scripted scenarios are enacted in sequence. The robot's performance in each scenario is qualitatively documented based on its navigation, worker-robot interaction, and support performances. Additionally, data such as the onboard computer's screen recording and LiDAR scans are captured at short intervals for diagnostic purposes[8]. Four time-lapse recording cameras

[7]Basically, the camera is specifically targeted for close-range, elevated, protruding, thin pipe-like structures, therefore scan points beyond 0.75m are discarded.

[8]This consideration is due to the continuity of worker operations and the robot's battery life. Data logging activities, including starting and stopping in actual operations, require research members to intervene and type on the robot's keyboard. These intervals can be quite interruptive for the demonstration, especially while the workers are deeply engaged in interacting with the robot. Therefore, although data logging of the entire demonstration process was intended prior to the demonstration, the group adapted contingently and respected the workflow.

are activated for recording prior to the demonstration. For each script, the workers use the controller to operate the robot for two runs. We kindly ask the workers to note feedback and observations during the demonstration, especially on occasions where they feel uncomfortable around the robot. Additionally, to increase the social navigation difficulty to a state closer to the initial carpentry formwork observations, group members occasionally obstruct the robot's path during navigation to test its adaptability and robustness when encountering humans. With no intention to cherry-pick the best runs, the following subsections document the robot's performance, both successes and issues, in each script, based on the review of time-lapse recordings. In the case of an issue, I offer a brief explanation of the potential cause based on my understanding of the technical framework, diagnostic logs, or similar issues previously encountered in the lab environment.

Script 1

Figure 5.7 illustrates a robot trajectory (one run) during the first scripted scenario. The robot successfully navigates from the current work zone to the central workbench, following the workers' commands. The robot's path is relatively smooth, with minor pauses and adjustments when encountering passages, workers, or group members. Notably, the robot chooses a slight detour in terms of global path planning, going around the material piles instead of through them. This is due to the contingent planning at the time of the navigation when participating workers actively walk through and around the material piles, a behavior relevant to the their work content. Given these circumstances, the robot's global path planning method opts for a safer but slower route to avoid further congesting the passages. The robot's navigation speed is moderate, neither too fast nor too slow, allowing it to keep pace with the workers without causing any disruptions, which is

an intentionally programmed feature. When a group member deliberately and suddenly obstructs its path, particularly around the central workbench, the robot promptly slows down or adjusts its direction to avoid collision or discomfort.

Script 2

Figure 5.8 illustrates a robot trajectory during the second scripted scenario. The robot promptly responds to the workers' remote commands, navigating from the central workbench to the designated work zones with loaded tools and materials. Similar to Script 1, the robot chooses a slight detour to avoid the material piles and workers. During the navigation course, group members intentionally cause more disruptions to the robot by initiating both abrupt front encounters and side crossings. The robot halts at one such occasion, a "freeze", and then slowly re-adjusts its direction at both a global and local planning level to avoid running into the group member. On other occasions, the directional adjustment around group members is relatively swift, leading to no apparent pausing or freezing. Upon arrival at the work zone, the robot manages to finish the delivery task without any major issues. However, the position it stops at is slightly distant from the worker duo (approx. 1.5m). The cause of this "drift" is due to the robot's localization. Since group members intentionally obstruct the robot multiple times in this script, and during those obstructions, the robot performs several rotations and direction adjustments in a short time. Compared to Script 1, the minor drifts in the turns and rotations accumulate more, eventually resulting in a slight mismatch between the robot's state estimation and its actual position[9].

[9]In simpler terms, the robot "thinks" it has arrived at the destination, but it, in fact, has not.

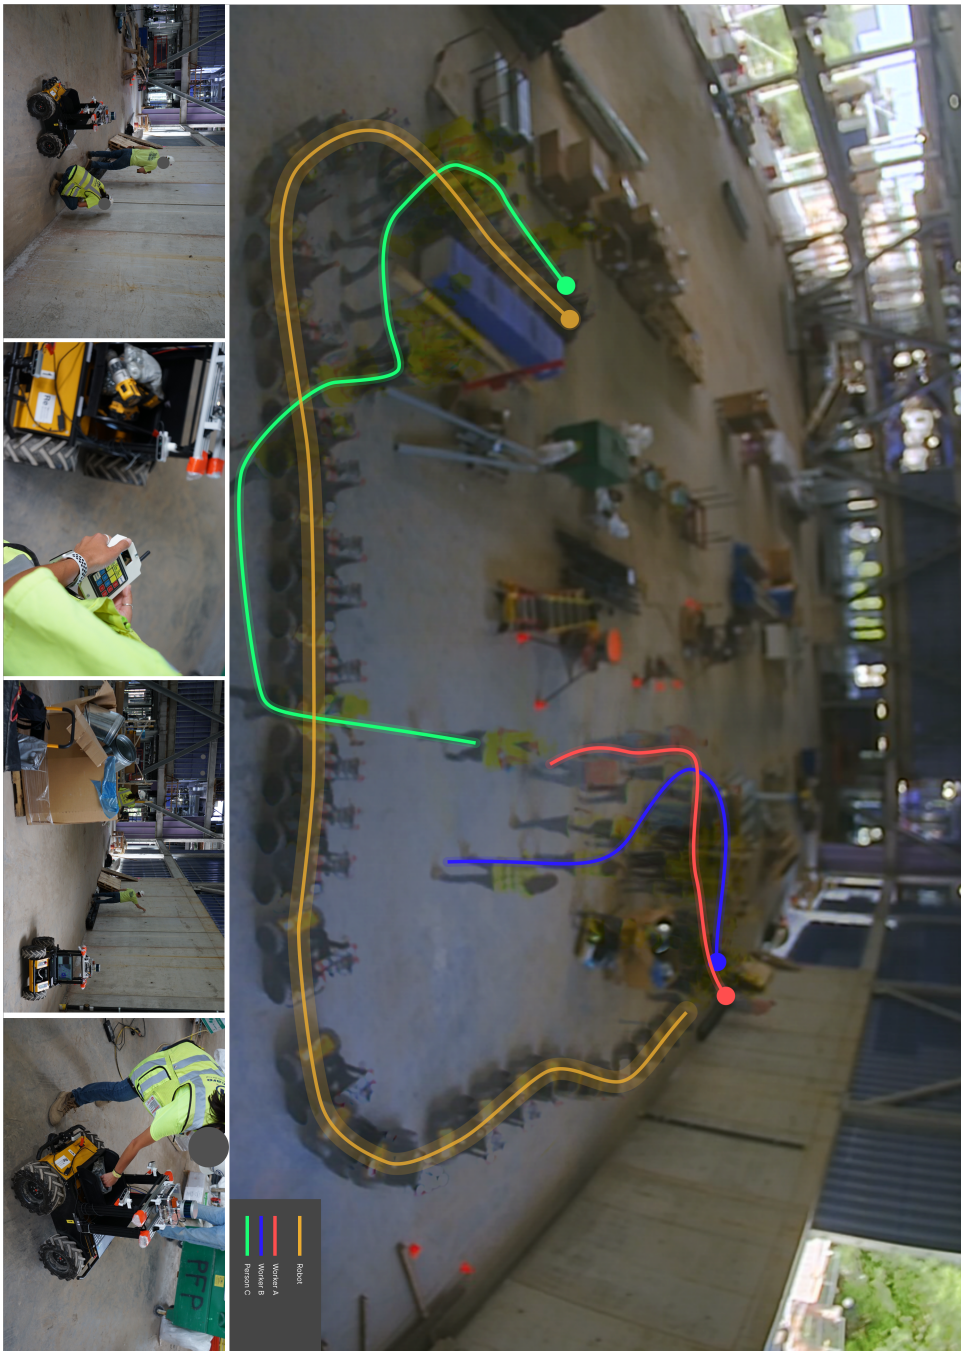

Figure 5.7: Robot trajectory from the time-lapse recording and photo documentation of Script 1. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a].

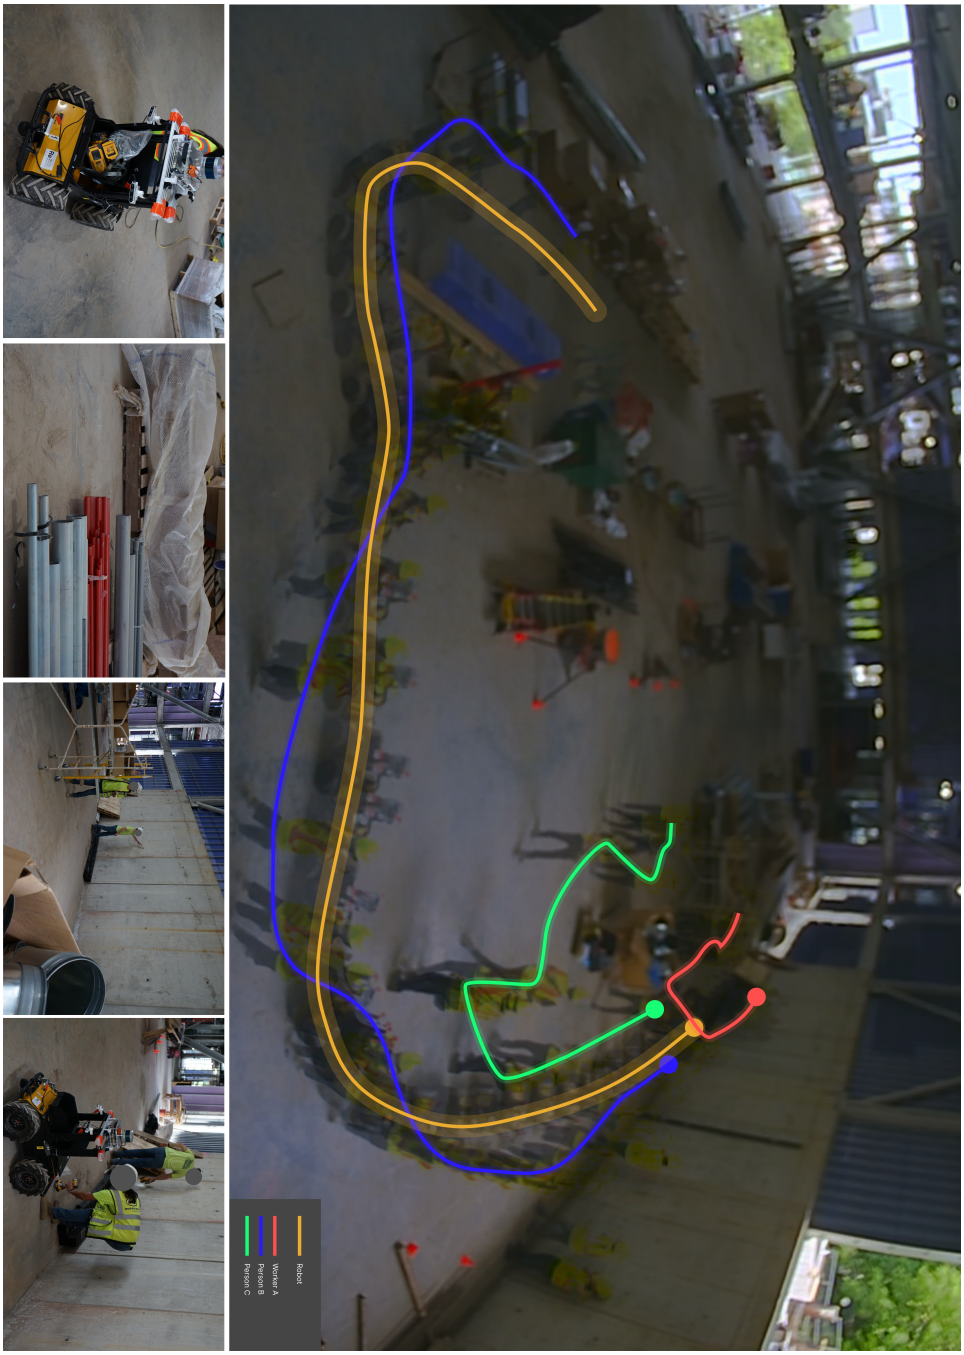

Figure 5.8: Robot trajectory from the time-lapse recording and photo documentation of Script 2. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a].

Script 3

Figure 5.9 illustrates a robot trajectory during the third scripted scenario. The robot successfully accompanies the workers from one work zone to the next, transporting tools and materials. The navigation is mostly smooth and adaptive, with only minor pauses and directional adjustments. In this scenario, tools and materials, such as drills and bags of metal plates are placed into the robot's container, which it effectively transports between work zones. The robot demonstrates its ability to adapt to the workers' movements, maintain a safe distance, and ensure social comfort. Compared to the first two scripts, this script presents a more congested operation space for the robot. Both the starting and goal work zones are located by the wall, limiting the robot's options for detours or drastically different routes. When navigating this corridor-like space, the robot frequently and actively adjusts its course based on real-time worker detection and tracking results, showcasing its ability to handle close encounters in tight spaces.

Additional Test: Long-Distance Navigation Challenge

The robot's performance in the additional test case, designed to challenge its navigational capabilities in sparsely mapped and unfamiliar areas, is less successful. The robot encounters difficulties in accurately localizing itself and navigating through the uncharted area. The lack of mapping and the accumulation of localization errors lead to an early, half-way termination of the navigation. The research group identifies this issue of contingent localization in dynamic and sparsely mapped environments as an intriguing future research direction (see Chapter 7).

Figure 5.9: Robot trajectory from the time-lapse recording and photo documentation of Script 3. Some photos were published in [Wu et al., 2024b] and [Wu et al., 2024a].

5.3.5 Post-Demonstration Worker Interview

Following the demonstration, a qualitative evaluation was conducted through a group interview with the participating workers, focusing on three main aspects of the robotic framework and its support functions. Much of this analysis has been previously reported in an earlier publication [Wu et al., 2024b], with additional details and findings elaborated here.

- **Robot navigation safety and comfort**: Workers were asked about their impressions of the robot's navigation behavior, speed, and safety when working alongside it, as well as overall comfort, including any potential disruptions caused by the robot. Most workers responded positively, affirming that the robot operated safely and smoothly within the work environment. They did not find the robot abrupt or intimidating and considered it a valuable addition to their workflow. Workers did not view the robot's involvement as an intrusion or a threat to their existing work practices but rather acknowledged the relief it provided from some of the physical burdens inherent in their tasks.

- **Robot hardware design and usability**: This section of the interview gathered feedback on the robot's physical design, ease of operation, and controller usability. Workers appreciated the intuitive design and the added safety features, such as foam-covered sharp edges. However, they identified areas for improvement, particularly in the robot's load capacity and its ability to transport bulky, heavy materials more efficiently. Workers suggested upgrading the controller to a more intuitive interface, possibly integrating smartwatches, which are commonly used in their daily routines. Regarding the tool/material container, workers found the size appropriate for the tools involved, and the retractable feature of the container improved accessibility.

- **Concept of robotic support within the workflow**: Workers discussed the robot's utility in supporting carpentry formwork tasks and offered suggestions for enhancing its role. They confirmed that the robot reduced on-site travel during formwork installation, streamlining certain tasks. A member of the research group elaborated on the rationale for selecting carpentry formwork as the focus, presenting site traversal data observed in earlier studies. Workers generally agreed with this focus and offered recommendations for further functionality. For example, they raised concerns about assuming a worker's presence at the central workbench and suggested implementing autonomous loading and smart inventory management to further extend the robot's capabilities.

The general attitude of the workers towards the robot's presence and support is positive, expressing a willingness to include the robot in the construction context and in their daily work. A notable observation from the interview is the workers' eagerness to extend upon, rather than entirely topple or reimagine, the current robotic framework towards more needed support functions and scenarios, which is a positive sign of their view of the robotic support. For instance, the notion of adding an arm to the robot to enable automatic loading and dexterous manipulation capabilities appeared multiple times in the conversation. Another suggested capability is whether the robot's current technical framework can be easily migrated to a larger robot that can carry stacks of wood boards or heavy metal crates. A third practical suggestion is to apply the robot to scenarios where workers urgently need to distribute small but numerous hardware items, such as screws, nails, and bolts, particularly during the interior installation and decoration phase, which follows the phase the group previously observed in the site study. These collective thoughts and suggestions from the workers

are valuable first-hand insights for the research group to further develop the robot's support capabilities and expand its applications in the future, providing a broader perspective on how and where the developed robotic framework can be applied and extended.

5.4 Quantitative Evaluation in the Lab

After the on-site demonstration, we conduct quantitative analyses and ablation studies in our laboratory at CMU Mill-19 Advanced Research Center. These studies aim to validate the robotic support, particularly assessing the use of RL for close-range navigation comfort around workers in key carpentry formwork scenarios. The hierarchical nature of our navigation stack allows us to compare the robot's performance with and without the RL-based social navigation layer, contrasting our RL-driven framework with plain efficiency-driven collision avoidance on the robot prototype. Additionally, we evaluate the impact of contextual fine-tuning and alignment in the RL model by comparing the performance of a vanilla RL model trained in generic settings with the fine-tuned RL model curriculum trained in work-specific scenarios.

5.4.1 Lab Sandbox Setup

To simulate on-site conditions and challenges in a controlled physical environment, we reconstruct a sandbox in the lab (Figure 5.10), capturing key geometric and spatial features of the construction site. The setup includes a central workbench and work zones, arranged to reflect the layout of carpentry formwork. Due to space constraints, the material piles are scaled down in volume but are intentionally placed and stacked disorderly to reproduce real-world clutter. While we cannot transport the exact objects

and materials from the site into the sandbox, we aim to gather objects with similar diversity and geometric characteristics. The objects include boxes, toolboxes, wood racks, thin wood beams, pipes, construction tools, metal hardware, and foam block clusters to mimic unstructured material piles[10] (Figure 5.10, Figure 5.12). We intentionally arrange occurrences of problematic cases observed on site; for instance, pipes and thin beams are placed horizontally and elevated to mimic the "extruding pipes" issue previously mentioned.

For ease of reference, we assign a 2D grid to denote different areas in the sandbox (left in Figure 5.10). The horizontal axes are labeled `A`, `B`, `C`, representing the space along the target wall, between the material piles, and near the central workbench, respectively. The vertical axes are labeled `1` to `4`, representing the spaces and passages from left to right. For example, the spot between the two left material piles is referred to as `B2`. This grid system helps denote the robot's and workers' positions during the evaluation.

5.4.2 Evaluation Scenarios and Test Cases

Two evaluation scenarios (Figure 5.11) were developed to correspond to the envisioned robot functions[11]. The first scenario concerns the *send* and *summon* functions, where the robot navigates between the central workbench (Row `C`) and work zones along the target wall (Row `A`) around workers amidst the material piles, through narrow passages. The second scenario focuses on the *accompany* function, where the robot accompanies workers in constrained spaces between work zones, primarily along the target wall

[10] Interestingly, most of these objects are collected from waste disposal from another construction site near the lab.

[11] In the on-site demo, there are three scripts and scenarios, while in the lab evaluation there are only two. This is because this evaluation primarily focuses on navigation performance and comfort. Other factors, such as distant controller operation and tool/hardware loading on the robot, are backgrounded. Thus, certain functions (i.e., send and summon) can be evaluated together in the same layout scenario.

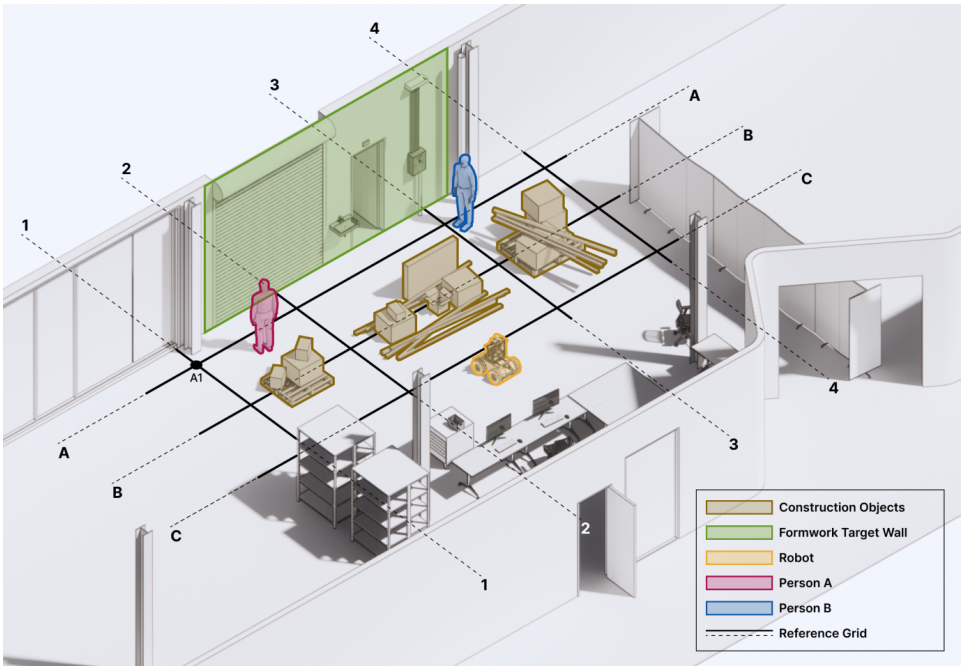

Figure 5.10: Laboratory evaluation setup. This image was included in [Wu et al., 2024a].

(Row A).

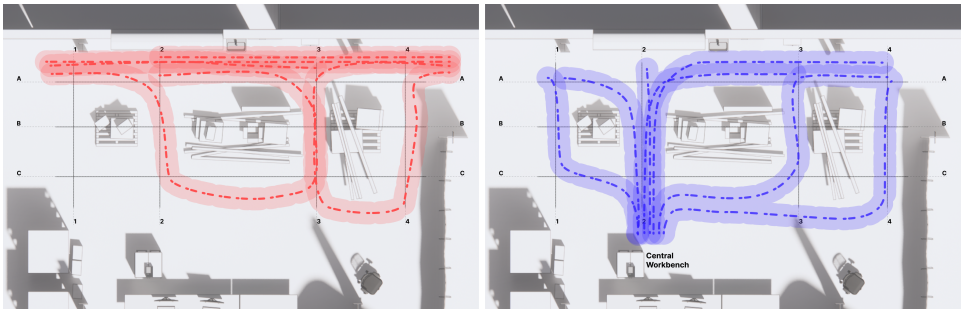

Figure 5.11: The two evaluation scenarios during quantitative evaluation: (left) scenario primarily tests the robot's navigation support capability among work zones by the target wall; (right) scenario focuses on the delivery capability between the central workbench and work zones.

In each scenario, 10 distinct test cases were developed to capture a wide range of worker activities and robot navigation objectives. This design is also briefly outlined in the earlier publication [Wu et al., 2024b]. Echoing the

studied worker duo, two research group members with prior knowledge of the carpentry formwork context participate as workers during the evaluation. In each test case, the robot's starting and goal positions, as well as each participant's starting position, goal position, and demonstrated work activities, are specified. Corresponding to the analysis in Chapter 3 and Chapter 4, participants primarily display three types of work activities: working beside walls, standing still, and walking around the site. The positions of these activities, such as the starting and goal positions of the walking, are specified using the grid system.

For each test case, 4 to 6 runs are conducted, blindly[12] alternating among non-RL[13], vanilla RL[14], and contextual RL[15]. During these runs, participants exhibit the same work activities, and the robot is tasked with the same navigation objectives. While it is infeasible to control human activities down to exact movements, it is generally agreed that participants start to display work activities when the robot is about 2.5m away and repeat their actions[16]. This approach aims to create similar worker-robot encounter scenarios across each run to relatively objectively evaluate the robot's performance. During each encounter and delivery, the participant closest to or directly engaged in the interaction serves as the primary evaluator to provide feedback and comments. The other participant observes, offers supplementary feedback, and documents the evaluation.

[12]This is to increase objectivity in the evaluation of worker-robot interaction

[13]RL-based social navigation layer is disabled, and the motion planning stack defaults to DWA

[14]Hierarchical motion planning with RL-based social navigation, but with a vanilla model without contextual alignment

[15]Hierarchical motion planning with RL-based social navigation, using the contextually aligned model

[16]Precisely controlling participants' activities in one test case across multiple runs is impractical as humans are not programmed robots. Using a clock to exactly time when to enact a specific action is also not constructive for investigating worker-robot interaction. Contingently, the robot's navigation process constantly varies. If participants' activities are displayed by time, then there might be runs when the robot arrives, but the participant's activity (such as walking) is already completed.

5.4.3 Evaluation Metrics

To evaluate the robot's performance with respect to its support capabilities, we use common metrics in social navigation, such as collision rate, success rate, robot freeze rate, and navigation time. Additionally and importantly, we develop two context-specific metrics to assess the robot's navigation comfort around workers: comfortable encounter rate and comfortable delivery rate [Francis et al., 2023]. A human-robot encounter is considered uncomfortable if the robot (1) gets too close to the worker (¡0.25m), (2) performs abrupt or drastic movements near workers that cause distraction or surprise, or (3) forces the worker to change their original path. Similarly, a delivery is considered uncomfortable if (1) the robot arrives at a position too far away for the worker to reach tools or materials from the container, (2) the robot performs too many unnecessary adjustments upon arrival, causing the worker to wait, or (3) the robot gets stuck during navigation and takes too long to arrive [Wu et al., 2024a].

5.4.4 Quantitative Evaluation Results

The results of the quantitative evaluation (Figure 5.12) are summarized in Table 5.1 and Table 5.2, which were also covered in [Wu et al., 2024b]. To visually present and compare the evaluation data between non-RL and RL based robot navigation, Figure 5.13 illustrates the trajectories of both the robot and participants, with notable occurrences during navigation marked. The results highlight the effectiveness of the RL-based social navigation layer in enhancing the robot's support capabilities in cluttered, worker-centric environments [Wu et al., 2024b]. In scenarios without RL, the comfortable encounter and delivery rates decrease significantly by approximately 0.4 and 0.15, respectively. While the implementation of RL does not significantly reduce major or minor collision rates (Table 5.2), it substantially

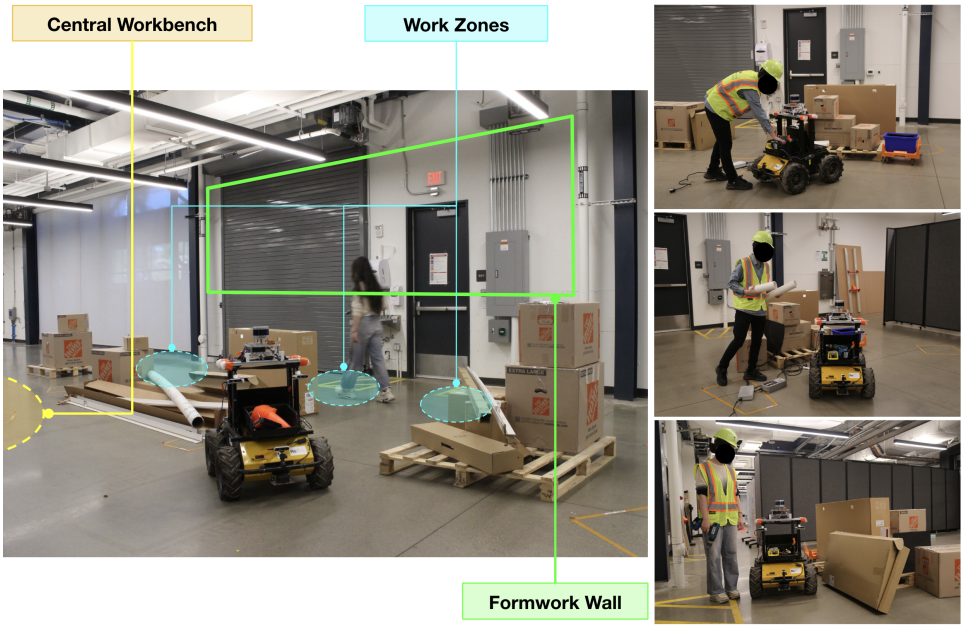

Figure 5.12: Photos captured during the lab evaluation. Some photos were published in [Wu et al., 2024b].

decreases obstructions to workers, robot freezes, and potentially dangerous movements. This leads to smoother, more socially aware, and comfortable support for workers, validating the initial design intentions.

	Comf. Delivery	Comf. Encounter
DWA	0.83	0.57
DWA+SARL (Vanilla)	0.95	0.87
DWA+SARL (Contextual)	**0.98**	**0.92**

Table 5.1: Quantitative results in lab evaluation regarding context-specific, comfort-related metrics [Wu et al., 2024b].

Qualitatively, this difference in navigation smoothness and comfort is also reflected in the participants' firsthand feedback and observations. Without the RL-based social navigation layer, participants often express discomfort and unease around the robot, particularly when approaching the robot from the front. The robot often comes to an immediate stop, then wiggles

	Success	Collision	Freeze	Time
DWA	0.83	0.05	0.95	29.40
DWA+SARL (Vanilla)	1.00	0.07	**0.23**	25.59
DWA+SARL (Contextual)	1.00	0.05	**0.10**	25.89

Table 5.2: Quantitative results in lab evaluation. "Success": successful delivery rate. "Collision": major/minor collision rate. "Freeze": robot freeze rate. "Time": navigation time (seconds) between two benchmarking work zones [Wu et al., 2024b].

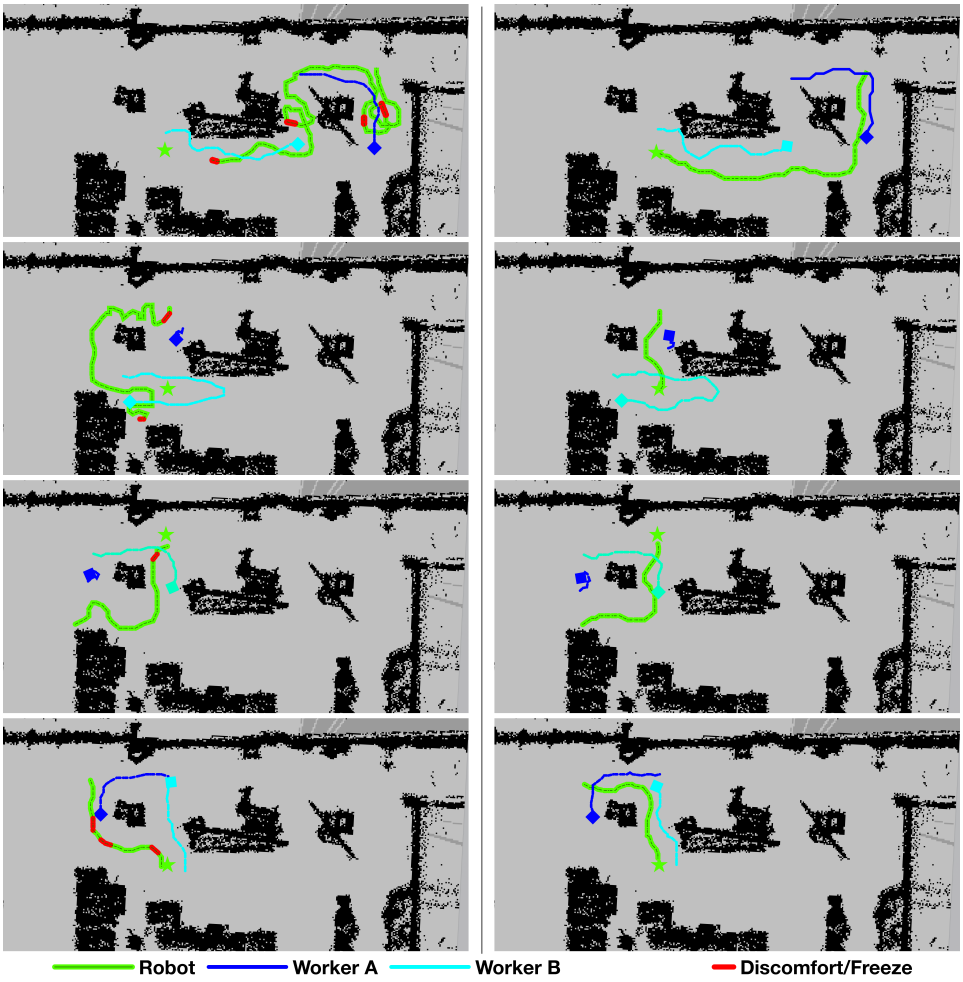

Figure 5.13: Trajectory comparison between non-RL (left) and RL (right) based robot navigation. This image result was published in [Wu et al., 2024b].

and rotates left and right without moving until the participant moves away. Such occurrences of robot "freeze", which have been frequently documented in other literature [Trautman and Krause, 2010], are reduced when the RL model is enacted. In observations, with the RL-based social navigation layer enacted, the robot usually starts to adjust its course in advance of the encounter, often a few steps away from the participants. This early adjustment is often achieved by turning a few degrees away from the participants' anticipated course of activity, leading to relatively smooth and adaptive robot trajectories. This robot behavior, also visually discernible in Figure 5.13, is less likely to cause human discomfort and matches the social awareness desired in the design of the robotic framework.

Additional analysis also reveals that the contextually aligned RL model, compared to the vanilla RL model trained only in generic settings, shows an improvement in the comfortable encounter rate and a decrease in the robot freeze rate, resulting in less obstruction or interruption when offering support. This enhancement is attributed to the additional alignment tailored to work-specific scenarios, which better equips the robot to handle unique patterns and attributes in our context [Wu et al., 2024b]. For instance, without additional alignment, the model is not fully prepared for cases when a participant repeatedly moves back-and-forth, leading to frequent and sudden 180-degree changes in direction. This is just one instance where alignment makes a difference. With similar prior training, the robot shows better adaptivity in handling these abrupt cases in advance, as numerically reflected in the mentioned metrics.

Chapter 6

Reflection and Discussion

6.1 Locating Robotic Technology Design: The Who, The What, and The How

One might notice the frequent mention of "context" and "contextualization" throughout this thesis. These terms capture a significant aspect of this research: the process of "locating technology design" [Suchman, 2002]. In her original text, Lucy Suchman highlights the importance of understanding technology within its specific social and organizational contexts, advocating for a design approach that accounts for the particularities of real-world practices. Inspired by such insight, this research emphasizes the need for specificity rather than abstract, visionary discussions of robotically-supported collaborative work. The goal is to ground the design and deployment of robotic systems in the unique contexts in which they are applied.

With the descriptions in Chapter 3, Chapter 4, and Chapter 5 in hand, this section reviews and reflects on three key aspects: "the who", "the what", and "the how", critically inspired by Suchman's vision. This reflection aims to uncover the underlying principles, positions, and high-level insights that

have driven the development of this robotic framework, highlighting how these principles are deeply integrated into the technology design process presented in the previous chapters.

6.1.1 The Who: Identifying the Users and Their Work Activities

In retrospect, this research did not initially set out to target carpentry formwork installation workers. This evolution in focus is why this thesis refers to "construction workers" in its title, rather than directly specifying "carpentry formwork installation workers". The project began with the rough, conceptual intention to develop robotic technology supportive of physically intense human labor, situated in the construction context. The identification of carpentry formwork installation workers as a key user group emerged from an iterative, "zooming in" process, which I would describe as "bidirectional alignment". In one direction, this alignment involved identifying specific challenges within construction labor where advances in AI/ML and social robot navigation could be effectively applied. As mentioned in Chapter 3, several other work practices were also reviewed and discussed within the group. Carpentry formwork was eventually chosen because the frequent site traversing burden observed in site studies is a pain point the current semi-/autonomous robot navigation techniques are suitable to resolve. In the other direction, the alignment required assessing the capabilities, especially the limitations, of available AI and robotics resources to find a concrete, realistic application scenario. While other work practices are also intriguing for research, some notably require further robotic manipulation dexterity or spatial flexibility, exceeding the group's current robotic resources and research capabilities. With no intention to pursue an omniscient "Frankenstein" robot, the aim of this research is to support a specific group of users

well, rather than projecting an all-encompassing, universal image of robotic assistance. This approach deliberately steers clear of grandiosity and unrealistic optimism about what the robot can currently achieve in the real world. However, this consciousness is not intended to confine the potential of technology or to suggest that AI/ML and robotics can only aid carpentry formwork installation workers. Rather, it strives for precision, effectiveness, and practicality in technology production, aiming to showcase the novelty of technical systems in a tangible, contextual, and meaningful way. This specificity also seeks to pave the way for the ultimate goal of integrating the technology into an actual workflow, where it can be solidly evaluated by workers in grounded, constructive conversations.

6.1.2 The What: Localizing the Specific Robotic Support Task

Following the decision to focus on carpentry formwork, the next critical matter then became: "What exactly is the robot supporting?". As detailed in Chapter 3, the research delves deep into the specifics of carpentry formwork through an analysis of work practices and the worksite. With these contextual understandings as a foundation, an important aspect of this research is determining the exact robotic support without arbitrarily assuming any drastic reconfiguration of the existing work. The current work practices, material and geometric features of the site, and worker behavioral patterns are studied and acknowledged to identify an immersive and naturally integrated way to introduce the robot. The three proposed support scenarios are all tailored to the specific links in the current workflows, not unrealistically dictating changes workers should make to accommodate the robot. Material and social challenges present in these tasks are not dismissed or viewed as mere obstacles. On the contrary, they are considered integral

to the support, serving as critical preliminaries for technology production. As consistently emphasized throughout the research, robots are positioned to support, not as replacements or disruptors, contrary to some laboratory research imaginaries.

6.1.3 The How: Realizing Robotic Support in a Contextually Informed Manner

With the "who" and the "what" identified, "the how" becomes a key element of the research to technically realize the robotic support. It particularly concerns how technology design is associated with, inspired by, and eventually situated in the uniqueness of the studied construction work. Two strategies are particularly notable in this regard: "unearthing human work through a technical lens" and "developing a deep tie between social and technical".

Seeing Human Work Through a Technical Lens

> "The way in which people work is not always apparent. Too often, assumptions are made as to how tasks are performed rather than unearthing the underlying work practices. By making the work visible, designers create a more intimate view of the workplace landscape."
>
> — Lucy Suchman, *Making Work Visible* [Suchman, 1995]

Inspired by the idea of making human labor visible [Suchman, 1995], this research aims to base technology design on an enriched understanding of real-world construction work, made visible through the qualitative methods employed by our group. Specifically, the focus is on carpentry formwork, in a manner that is both socially detailed and technically insightful. As

described in Chapter 3, this is achieved through the decomposition of both the worksite and workers' activities. The goal of this decomposition is not only to gain high-level qualitative insights but also to translate these insights into a meaningful dialogue with existing research in social robot navigation, machine learning, and human-robot interaction.

Developing this dialogue is not a one-time event, but rather an iterative "trial-and-error" process that evolves and refines over the course of the research. Key elements of the decomposition are gradually discovered and often initially considered as "interesting" factors during thought experiments, group discussions, literature reviews, and daily research activities. For example, the idea to closely investigate the behavioral differences between a carpentry worker and a common pedestrian emerged one afternoon. After spending considerable time training the RL algorithm in simulation, using symbolic, monolithic agents (circles) moving in simple paths, I was exposed to a raw video recording from the site, capturing the lively activities of the workers. The vivid contrast between these real-world activities and the simulation led me to intuitively question the simulation: "This doesn't resemble the tests I'm running at all". Delving deeper into this difference resulted in a more detailed analysis of worker activities, which ultimately led to the development of more diverse and generative agent behavior in training the RL algorithm, as presented in Chapter 4 and Chapter 5.

Another idea to investigate workers' activities through both an individual and collaborative lens comes from my colleague's persistence. He insisted on describing the worker duo's activities separately. This insistence started with using "Worker A" and "Worker B" to separately annotate and analyze the different roles of the main installation worker and the assistant. Seemingly trivial at first glance in a technical sense, it eventually turned out to be a critical factor in outlining the social dynamics of the work. If under-

standing individual activities helps diversify agent behavior in RL training, then understanding the collaborative dynamics guides the design of experiment scenarios for the algorithm and the prototype, which in turn critically captures the essence of the robotic support and worker-robot interaction.

The fragmented ideas and factors go through several iterations before synthesizing into a coherent, descriptive framework. Rather than a direct translation or a plain amalgamation, the synthesis is highly selective and prioritized towards the specific robot support. For instance, while qualitatively differentiating the work zones helps understanding the spatial progression of the work, it is less relevant to do so when resolving the robot's support and navigation capabilities. This is why all zones are currently reviewed generically, and are noted only as A, B, and C, not with the analytical depth and attention placed on worker activities or tools/materials. As a reflection, I feel such prioritization is necessary. While a spanning coverage of the qualitative analysis is important, the depth and relevance of the analysis towards the intended technology design is even more critical.

Develop a Deep Tie Between Social and Technical

Besides explicitly developing a framework to decompose the work, the "how" is also achieved through forming a deep tie between ethnographic understanding of the work and the technology design process. As shown in multiple cases in Chapter 3 and Chapter 4, the tie is infused in many links of the framework. From the formation of the robotic hardware to the development of the RL algorithm, to the design of the worker detection and tracking module, to the simulation and real-world testing, multiple technical decisions are shaped by the social and material features of the work. As a general principle, technology is not selected merely as a pursuit of fancy and state-of-the-art but also weighs in suitability within the specific context.

Taking the worker detection and tracking module as an example, with the fast development of computer vision and autonomous driving techniques, new methods appear on a monthly, if not weekly, basis. However, recognizing the inherent differences between a highly complex construction setting and what is implicitly assumed in those emerging research, and especially how these differences may affect algorithmic performance, helps maintain a sense of consciousness in the prototyping, particularly within a limited time frame and with limited resources. This sense of technical suitability, seemingly empirical and personal, is largely fueled by the deep understanding of the social and material gist of the work gained through immersive observation and discussion in the research group.

Besides the direct impact on technical decisions, the tie also triggers and inspires theoretical thinking in technical research. For instance, addressing the behavioral differences between carpentry workers and pedestrians, at a conceptual level, concerns how to effectively fine-tune a generically trained RL model to align with specific preferences and tendencies. Combined with certain expectations such as model performance consistency in Sim2Real, which emerged as an actual challenge during the robot's development, inspired the idea of the "contextual fine-tuning of RL" module, as described in Chapter 4. The module itself is an algorithmic pipeline that is generalizable to other robotic scenarios, not just limited to carpentry formwork or construction robotics. But it is this specific context, especially its inherent challenges and constraints, that nudged its development. As a reflection of such inspiration, I would say that the tie between social and technical is also bidirectional. The social understanding of human work helps situate and shape the technology design in concrete ways, while also serving as fertile ground for formulating and thinking about technical problems that might otherwise go unnoticed or be absent in a generic, abstract, and detached research setting.

6.2 Architecting A Robot For Real-World Challenges

Construction sites resemble some of the most complex, populated, and dynamic environments for robotic applications. Orchestrating a technical architecture for a "work companion robot" capable of operating in such challenging contexts, particularly with the ultimate aim of tangibly supporting existing workflows and workers, is a difficult task. In this section, I reflect on the "architecture" process of this robotic framework. In some cases, I also wish to draw parallels with architectural design principles and practices that I find useful, particularly those that can be meaningfully applied to the robot design and prototyping process.

The first insight is about balancing the high-level vision with the low-level technical details. The aim of this research——robotically-supported collaborative work——requires several robotic capabilities, each with its own technical complexities. For instance, the robot must perceive and locate itself in a dynamic environment, navigate through clutter, detect and track workers, and adaptively support them with social awareness. Each of these features could be a research topic on its own, but the robotic framework integrates them to serve the ultimate purpose of robotic assistance. Take the worker detection and tracking module, for example. It involves multiple sensors, integrating detection and tracking methods, and optimizing inference speed. Yet, this module is a means to an end——supporting the robot's socially aware navigation. Maintaining focus on the broader goal, without getting bogged down in subproblems, is critical. Even the contextual fine-tuning of the RL module, though algorithmically complex, is developed with the clear aim of aligning the robot's navigation with workers' bespoke behavioral patterns. This approach——balancing technical details with the overarching goal of robotic assistance——can be adapted to other human-

centric contexts beyond construction. The same principles of grounding high-level vision in technical capabilities apply across fields where robotics integrates into specific work environments. Maintaining this focus has been crucial for building and managing complex technical systems.

The second insight is about embracing an imperfect, candid view of the real world. Developing robotic prototypes in the lab versus in a real-world setting involves different conditions, premises, and difficulties. These often underestimated differences, if not acknowledged or articulated, may become decisive pitfalls in the actual efficacy of the robot. This concern is particularly rooted in my experience as a former architect, who has been on construction sites, witnessing how some lab-incubated robots and machines may demonstrate well but eventually perform poorly in real-world scenarios. In this research, viewing the context—a cluttered, unstructured, and dynamic construction site as it is, rather than in a simplified, idealized, or even detached fashion—is an important angle I aim to maintain. A great deal of the effort in Chapter 4 unfolds and revises itself upon these somehow tricky, seemingly trivial but bold realities. For instance, the robot's sensor package is repetitively upgraded to address corner cases observed on site. Without exposure to an actual construction site, such dedication can easily be dismissed in the lab. Yet, these have shown to be the very factors that can make or break the robot's performance in the real world.

The third insight relates to the iterative, clumsy, rough, and bumpy process of technical development. A typical architectural design process involves three stages: coarse design, fine drafting, and final detailing. This methodical deepening is mirrored in the prototyping of this research's technical architecture. One may notice that in Chapter 4, such iterations often start with the simplest, most broadly applicable, even somewhat "outdated" solutions. For instance, the research's initial trial of the navigation stack was

the most basic ROS navigation stack. This decision is not due to a lack of knowledge but is based on the principle of establishing a functional baseline before pursuing optimization and customization. It also underscores that architecting a robot is not solely about chasing the latest numeric benchmarks. Principles such as time-tested reliability, contextual appropriateness, and engineering pragmatism still matter. Therefore, recent methods are not pursued baselessly, but with grounding in comparisons with more established approaches, which is showcased in the "hierarchical motion planning" module.

Bearing these insights in mind, we can see the multifaceted, complex, and iterative nature of developing a robotic architecture for real-world challenges. The process is not just about technical prowess or theoretical sophistication. It is also about the systematic structuring, building, and handling of solutions responsive to concrete realities and complexities.

6.3 "Hardware is Hard"

The intricacies of handling a real robot prototype extend far beyond theoretical models and simulations, diving deep into the tangible world of hardware. This difficulty, often referred to as "hardware is hard" within the research group, encompasses a multitude of challenges, from dealing with long-tailed issues to mastering the minutiae of the robot's nuts and bolts. This section reflects on these complexities and offer some high-level experiences in overcoming the hardware difficulties in this research.

The first reflection relates to the handling of long-tailed issues. While interacting with real robots, one quickly learns that the diversity of potential issues forms a long tail—rare and unexpected problems that are difficult to predict and often even harder to trace. Solving these long-tailed issues

requires a deep understanding of the hardware at both a detailed and systemic level, enabling one to draw intuitive connections between seemingly unrelated factors. For instance, the utilization of a GPU is essential in this research. However, having a GPU running on edge with multiple deep neural network inferences can result in a hot battery. Initially, without explicit heat management, the group observed that the robot would occasionally stop functioning with errors like "GPU falling out of the bus", leading to discontinuous, jagged navigation performance. Despite the software being well-coded and tested, the issue persisted. After several weeks of debugging, the problem was pinpointed to the keyboard's flashy backlight, which, on our specific computer model, made the upper surface and battery extra hot, causing the GPU to overheat and disconnect. The solution was simple: turn off the keyboard backlight and monitor the computer's heat. This example vividly illustrates how seemingly trivial hardware factors can have a profound impact on a robot's overall success. Two principles are critical when dealing with these issues: maintaining a broad perspective, even when dealing with minute details, and always paying attention to hardware basics. Even with well-developed, industrial-grade robots, it is crucial to monitor aspects like battery, heat, power management, and cable connections.

The second lesson highlights the profound difference between theoretical knowledge and practical implementation. Applying advanced algorithms or methods to real robot hardware is never a simple "plug-and-play" process. The performance of a promising algorithm in simulation or on a dataset does not guarantee its success on one's particular hardware suite. For example, the implementation of LIO-SAM for site mapping, despite the original author open-sourcing a detailed repository of the algorithms on GitHub, required significant effort. The algorithm relies on a tight coupling between 3D LiDAR and IMU, necessitating a deep understanding of the robot's sensor suite and the algorithm itself. The physical setup of our modular LiDAR

box, as mentioned in Chapter 4, needed to be accurately represented using a few parameters and matrices. Incorrect parameters or misunderstandings of the algorithm's assumptions can lead to issues such as constant drifting of the registered point cloud or misalignment of axes and transforms. Resolving these issues required a long-winding, hands-on, trial-and-error process. Furthermore, our IMU model posed additional challenges for the method's initialization. For the algorithm to work, the data stream of the IMU and 3D LiDAR must appear in a specific order. This seemingly trivial requirement became a decisive painpoint on our hardware. After multiple software-level attempts to maintain such an order, the group ultimately gave up and developed a hardware startup sequence instead to circumvent the issue. Two critical experiences emerged: do not underestimate the engineering gap between a theoretical method and its application on the hardware, and be prepared to adapt and iterate as necessary based on actual performance.

The third lesson pertains to the system management of the hardware, focusing on sharing tacit knowledge and establishing operation protocols. In this research, many brilliant master's students[1] contributed to the experimentation of the robot, accelerating the group's exposure to various hardware issues and helping to find solutions. Instead of passing these tacit, hands-on experiences casually or logging them in personal notes and dev logs, I decided to establish a formal, systematic knowledge base for hardware management and documentation. This sharing facility includes a detailed list of common hardware setups and configurations, issues and their solutions, and preventative measures, operation manuals and protocols, serving as a repository of collective wisdom for the research group. This facility, though not included in the research publication, has proven invaluable in managing and driving the robot prototype and facilitating the group's collective learning.

[1] Please refer to the Acknowledgement

Fourthly, embrace simplicity. Many hardware solutions in making the robot work are surprisingly simple, despite the complex journey to discover them. The phrase "dummy things that make robot work" encapsulates this notion. However, the value of these solutions lies not in their complexity but in their efficacy. Finding these straightforward yet effective solutions is a significant aspect of the research, exemplifying the creativity and persistence required to navigate the multifaceted challenges of working with hardware. For instance, the worker detection and tracking module's "tilted camera" solution emerged after realizing that workers' faces would disappear from the camera's FOV, influencing the detection accuracy. Initially, the most promising solution seemed to be adding another camera on top of the robot, which was costly and time-consuming. After brainstorming and testing on a smaller robot, the simpler solution of tilting the camera up was implemented, effectively resolving the issue with minimal effort.

Finally, every bug encountered and every problem solved contribute to a growing body of knowledge. This accumulation of experiences and familiarity with various issues gradually transforms a novice into an adept researcher.

6.4 Is Reinforcement Learning the Answer?

This research demonstrates a full-stack exploration of reinforcement learning (RL), both theoretical and applicational, through a companion robot adaptively supporting construction workers, almost like opening up a Pandora's box of potentials and pitfalls. This section reflects on the advantages and disadvantages of employing RL as an adaptive strategy for robotic applications, particularly in navigating the complex and nuanced dynamics of human workspaces, in hopes of shedding some high-level insights into

leveraging this machine learning paradigm.

One of the most significant challenges in applying RL to solve real-world robotic problems is its voracious appetite for data. RL algorithms, especially those based on deep learning, require vast amounts of experience to learn even moderately complex behaviors. This hunger for data presents a practical challenge especially in scenarios highly associated with human activities such as construction work, where capturing and collecting actual human activities and interactions can be time-consuming, expensive, and sometimes impractical. The reliance on simulated environments, as described in Chapter 4, becomes a double-edged sword. While simulations offer a controlled, scalable way to generate data, the gap between simulated and real-world experiences—also known as the "sim-to-real" gap—poses a substantial hurdle. In this research, such a hurdle is exemplified by the frequently mentioned difference between a formulaic walking agent in prevailing social navigation simulation environments and the actual construction worker. The challenge of bridging this gap is not just about the fidelity of the simulation but also about the nuanced, context-specific behaviors that are difficult to reproduce in a simulated environment. This challenge is further compounded by the need for continual fine-tuning and adaptation, as the robot encounters new scenarios and interactions in the real world. Without the foundational efforts of the community in developing and refining simulation environments, the path toward applying RL in the real world would be steep and fraught with uncertainty, especially for scenarios with intricate social dynamics. In fact, if there were no previous efforts by Jean's group and other research groups on developing open-sourced, customizable social navigation simulation environments, the research would probably reconsider whether to employ RL in the first place.

RL's strength lies in its ability to master skills through extensive practice

and trial-and-error, a learning paradigm that mirrors some aspects of human skill acquisition. This capability makes RL well-suited for tasks requiring a degree of delicacy and adaptability, aligning closely with certain human approaches to learning. However, this is not the only way humans learn. Frankly speaking, this trial-and-error honing of an atomic, reward-driven skill, such as how to twist a pen with fingers, often appears when there are already some foundational physical capabilities in place as premise. RL's focus on outcome-oriented learning, without an explicit reasoning or implication of the task mechanism, introduces inherent disadvantages. This is particularly concerning for contexts like robot navigation in complex workspace. Unlike mindlessly dodging awkward encounters when walking on the street, even humans need to keenly observe the surroundings and navigate based on contingent information: "the path is clear because there is no ongoing work", "I should take a detour because I should not be crossing a busy worker duo". These additional "situated awareness"[Suchman, 2015] essential to survive an environment such as construction site, compared to other simpler, more intuitive scenarios, are currently not present in the formulation of RL solutions. One may argue that if we include all these factors into the RL's state space and make the model large, the problem can still be solved. But at what cost? Moving towards that end, RL might not be the most suitable solution.

Despite its noted disadvantages, RL demonstrates a unique strength in refining nuanced responses in bespoke situations, particularly when aligning a generically pre-trained model that has already acquired baseline capabilities. This strength mirrors the training paradigms of large language models (LLMs), where reinforcement learning is used to align models with human preferences. The successful adaptation of RL models to the specific carpentry formwork scenario, a technical highlight of this research, underscores the potential of RL. However, I contend that RL should not be viewed as a

standalone, one-size-fits-all solution, but rather as a component in a more integrated approach to AI-driven interactive robots.

The future of RL in robotics, especially in socially and physically complex scenarios, likely lies in a hybrid approach that combines RL's contextual fine-tuning capabilities with the understanding and reasoning provided by more foundational models. This integrated approach could offer a more comprehensive solution to the challenges of current learning paradigms, enabling robots not only to adapt through extensive experience but also to reason deeply about the social and physical environments in which they operate.

Chapter 7

Conclusion and Future Work

7.1 Contributions

This interdisciplinary research integrates computational design, robotics, and AI/ML to address the unique challenges of robotically supprting physically demanding human work, with a focus on the specific complexities of on-site manual construction. Choosing carpentry formwork as a prototypical scenario, a human-centered, AI/ML-driven robotic framework is developed, leveraging reinforcement learning methods to adapt a pretrained social navigation model to the observed work patterns and behaviors of construction workers, aiming to provide tangible support and companionship within workers' existing workflows. While this research focuses on construction, the methodology——basing technological designs on deep, qualitative understandings of specific practices——can be extended to other human-centric contexts where technology must align with domain-specific features, needs, and activities. The key contributions of this research are threefold.

First, this research contributes to the field of computational design by exploring an area less examined by previous studies. While previous re-

search primarily approaches the problem of constructing physical spaces from material-centered, design-centered, or robot-centered perspectives, this research directly tackles the substantial presence of human labor and the inherent complexities and "messiness" of work scenarios. By proposing a human-centered robotic architecture based on the actual work context of carpentry workers, this study addresses critical, often overlooked real-world challenges in meaningful ways. It also presents a case study of a technology design process developed in close dialogue with the social and material specificities of human work.

Second, the research contributes to the robotics domain, particularly the growing field of human-robot interaction (HRI). The research showcases a potential prospect of how social robots can be designed and introduced into existing, laborious work scenarios while offering situated and handy assistance, in addition to the current common image of public guiding, museum touring, food delivery, etc. The research also provides a solid and reusable technical foundation and platform to trigger more contextualized, empathetic, and considerate research in similarly physically-intensive work scenarios, in construction, manufacturing, agriculture, and even in space exploration, towards a robotically-supported future of work.

Third, this research contributes to the field of AI and machine learning, particularly in reinforcement learning (RL). It demonstrates how recent advancements in deep reinforcement learning can be applied to adaptively navigate robots in complex, human-centric work environments. The research also introduces a methodological approach for improving, aligning, and contextualizing a generic RL model to practical complexities and worksite practices, while addressing how real-world "uncertainty", "messiness", and "imperfection" can influence algorithm performance——factors often underexplored in typical algorithm benchmarking settings. This approach integrates insights

from social and qualitative research, along with consideration of Sim2Real transfer and robotic performance stability, achieved through both experimental validation and algorithmic optimization.

In more technical detail, the research contributes [Wu et al., 2024b]:

1. A robot prototype demonstrating on actual construction sites how reinforcement learning algorithms can work in tandem with supplementary methods to enable assistive robot tool/material deliveries and comfortable worker-robot interactions, central to supporting the daily work of construction workers.
    (a) An effective pipeline for employing, contextualizing, and fine-tuning a group of recent RL-based social navigation algorithms to enhance human-robot interaction in specific work contexts.
    (b) A modular and open-source design and implementation of a robot system and hardware, complete with a stack of control and sensory packages tailored to cluttered indoor work environments.
    (c) A robust architecture of perception modules designed for complex, unstructured obstacles and the efficient detection of workers' activities.
    (d) A lightweight and hierarchical motion planning stack that harnesses advantages of both RL-based socially aware path planning and search-based path planning methods.
2. A generalizable, efficient RL-driven robotic framework, encompassing both hardware and software, that serves as a navigation base for enabling flexible and mobile robot support functions in indoor, cluttered, labor-intensive work scenarios.
3. A concrete robotic foundation for future investigation into algorithmic, ethical, social, and applicational aspects of robotically-supported

collaborative work.

At a higher level, I hope this research advocates for and substantiates the vision of "human-centered AI and automation". I envision a future where robots empathetically support and work alongside people, akin to TARS in "Interstellar", MB-26 Benben in "The Wandering Earth II", and Baymax in "Big Hero". Coming from an architecture background and witnessing the long-standing, yet often overlooked, difficulties in heavy-manual traditional industries, I hope this research prompts reflection on the aims and visions of developing AI and robotics technologies in the first place. While the current trend in AI and robotics research is predominantly driven by the pursuit of efficiency, performance, and autonomy, I wish to remind the community of the importance of human-centeredness, contextuality, and empathy in the development of these technologies.

In addition to these contributions, I also wish to candidly share the tacit knowledge, difficulties, failures, and successes encountered during the research process through this thesis, with the aim of prompting deeper reflection on designing and building AI-driven robots for real-world challenges within the broader community. As the current norm, most published papers often document only the final technical choices-what equipment was used, what algorithm was selected, and what procedures were performed. While such clarity and brevity are essential for research publication, they largely omit the actual, in-context process of solving real-world problems, which inevitably involves iterations of twists and turns, trials and errors. As readers may find throughout this thesis, many seemingly trivial issues can profoundly affect the success or failure of the overall robotic framework. I resolved most of these issues through painful, iterative problem-solving, by watching numerous YouTube tutorials, reading countless blog posts, and scrutinizing pages of GitHub Issues. At that time, I often hoped for a more

explicit way of sharing these tacit knowledge and experiences. Therefore, in my doctoral dissertation, I hope the systematic and reflective documentation of this research, including both successes and pitfalls, can benefit future researchers facing similar difficulties, and the community at large.

7.2 Limitations

Despite the promising prospects, this research has several limitations worth mentioning. First, the current robotic framework does not include a robotic manipulation component. The research at this stage focuses on the navigation and mobility aspects of the robot, excluding the capability to manipulate objects or tools. Future work could explore the integration of robotic manipulation capabilities into the existing platform, potentially enabling more diverse support functions. For instance, equipping the robot with an arm could allow it to pick up and deliver tools and materials directly to workers, or engage in more contact-rich assistive manipulation tasks. The absence of this capability currently assumes that a worker is needed to load and unload the robot, which detracts from the overall efficiency of the system.

The second limitation pertains to the robot's navigation capabilities. While the research addresses many encountered issues and corner cases, certain challenges remain unresolved. As mentioned in previous chapters, localization drift is still present in the robot's navigation system, causing cumulative deviations from the intended path over time. This issue primarily arises from the complexity and non-stationary nature of the construction work environment. Future work could explore more advanced mapping and localization techniques, particularly those incorporating semantic-matching alongside scan-matching. For instance, static structures such as columns and

finished walls could be used as anchors, while temporary, constantly changing objects like material piles could be omitted. This selective approach mimics human behavior in similar situations and could leverage advancements in vision-based reasoning and SLAM. Utilizing the modular nature of the framework, future research can flexibly replace less performant modules with more advanced methods emerging in the fields of robotics and AI, aiming to achieve better overall navigation performance.

The third limitation relates to the demonstration site. As mentioned in Chapter 5, while the research group made its best effort to find a site representative of carpentry formwork conditions, it still differs slightly from the conditions observed in the preliminary site study. Although cluttered with various construction objects, the demonstration space is larger than the initially studied site, and the target formwork wall is built with fewer exposed rebars. Additionally, the evaluation is limited by the fact that it did not occur during an actual workflow, but was primarily conducted to document on-site, in-context reactions from actual construction workers. While each construction site is unique, it remains an intriguing question how these particular differences may impact the robot's performance and the overall system's usability. Future work could explore deploying the robot framework in environments that better match these specific conditions, if feasible, to further evaluate and optimize the framework.

7.3 Future Work

Moving forward, there are several high-level research directions for future work that can build upon the research presented in this thesis. Particularly, there are three directions that I believe are most promising and impactful, and I am currently pursuing them as promising research career goals.

The first direction, emerging from the insights from Chapter 6, concerns the deeper "reasoning" capability of robots within work scenarios, particularly regarding the interpretation of work ethics and norms. If we consider RL as the robot's control layer enabling the robot to act contingently upon the cumulative understanding of human activity patterns, then the "reasoning" layer needs to be developed on top of this layer to interpret the underlying reasons, associations, and intentions behind these activities. The fast-evolving fields of AI and robotics offer emerging methods and tools to explore this possibility. Examples include leveraging large language models (LLMs) as few-shot learners and interpreters of intuitive verbal instructions, using vision-language models to facilitate multi-modal context understanding, or utilizing knowledge graphs to represent and reason about the implicit knowledge and norms in human work. This direction could potentially lead to more context-aware robots that can better understand and support humans.

The second direction concerns the broader applicability of the robotic framework developed in this research. While the research focuses on the specific scenario of carpentry formwork, the technical framework itself is designed to be modular, expandable, and generalizable, and can be adapted to other labor-intensive work scenarios where robot support and companionship is beneficial. Future research could explore deploying the framework in other contexts, even beyond construction, such as manufacturing, agriculture, or space exploration, to expand and evaluate its performance and usability in different environments. This exploration could also involve further optimizing and evolving the framework's components, such as the perception modules, motion planning stack, hardware design, and machine learning algorithms, to better suit the specific requirements of these contexts. By doing so, the research could contribute to a more comprehensive understanding of how robots can support humans in various physically intense

work scenarios, and further substantiate the vision of robotically-supported collaborative work.

The third direction concerns the ethical and social implications of deploying robots in human work scenarios, especially how such support should be more feasibly evaluated before deployment. Rather than discussing these ethical considerations in a general, verbal sense, a future direction of this research explores how to embed these considerations into the design and development of a benchmarking utility, potentially a photorealistic, physics-based simulation environment, that can capture both the material and social dynamics of specific real-world work scenarios. In contrast to the symbolic simulator and a real-world sandbox, such a simulation utility can provide a safe and controlled space, a "digital twin" middle layer, for evaluating the performance and usability of robot systems (Figure 7.1) before their real-world deployment. This would create a space for investigating the criteria, ethics, and policies for robotically-supported collaborative work at large. Similar to the approach presented in this thesis, this future direction is inherently interdisciplinary. It could involve collaboration with experts in ethics, sociology, and human factors to identify and define key ethical and social considerations. It can also consider working with experts in computer graphics, animation, and the film industry to capture, gather, and even generate more realistic human activities in simulation. By doing so, the research could contribute to a more systematic, safe, and rigorous evaluation of robots in human work scenarios, ensuring that they are designed and deployed in ways that are ethical, accountable, and beneficial to human workers.

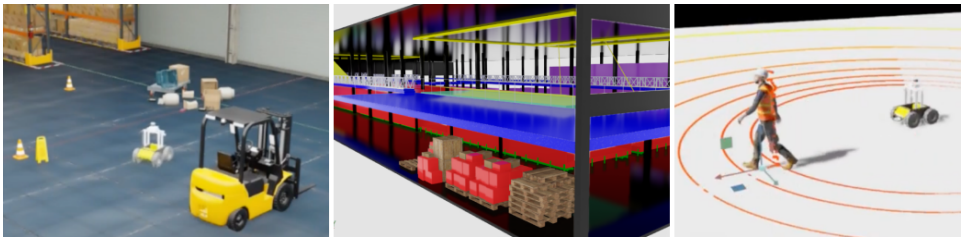

Figure 7.1: A preview and current progress of developing a photorealistic, physics-based simulation environment for benchmarking context-awareness of robotic system in work environment. Image credit: ReAC Research Group, 2023.

Bibliography

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.

Evan Ackerman. A robot for the worst job in the warehouse: Boston dynamics' stretch can move 800 heavy boxes per hour. *Ieee Spectrum*, 59 (1):50–51, 2022.

Sean Ahlquist. Integrating differentiated knit logics and pre-stress in textile hybrid structures. In *Modelling Behaviour: Design Modelling Symposium 2015*, pages 101–111. Springer, 2015.

Sean Ahlquist and Achim Menges. Frameworks for computational design of textile micro-architectures and material behavior in forming complex force-active structures. In *ACADIA*, volume 13, pages 281–292, 2013.

Sean Ahlquist, Julian Lienhard, Jan Knippers, and Achim Menges. Exploring materials reciprocities for textile-hybrid systems as spatial structures. In Michael Stacey, editor, *Prototyping Architecture: The Conference Papers*, pages 187–210. Riverside Architectural Press, Toronto, 2013. ISBN 978-1-926724-37-9.

Naveed Anwar, Muhammad Amir Izhar, and Fawad Ahmed Najam. Construction monitoring and reporting using drones and unmanned aerial

vehicles (UAVs). In *The Tenth International Conference on Construction in the 21st Century (CITC-10)*, volume 8, pages 2–4, 2018.

Aleksandra Anna Apolinarska, Matteo Pacher, Hui Li, Nicholas Cote, Rafael Pastrana, Fabio Gramazio, and Matthias Kohler. Robotic assembly of timber joints using reinforcement learning. *Automation in Construction*, 125:103569, May 2021. ISSN 0926-5805. doi: 10.1016/j.autcon.2021.103569.

M Sanjeev Arulampalam, Simon Maskell, Neil Gordon, and Tim Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on signal processing*, 50(2):174–188, 2002. doi: 10.1109/78.978374.

Reem Ashour, Tarek Taha, Fahad Mohamed, Eman Hableel, Yasmeen Abu Kheil, Malak Elsalamouny, Maha Kadadha, Kasturi Rangan, Jorge Dias, Lakmal Seneviratne, et al. Site inspection drone: A solution for inspecting and regulating construction sites. In *2016 IEEE 59th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 1–4. IEEE, 2016. doi: 10.1109/MWSCAS.2016.7870116.

Ahmed Attalla, Omar Attalla, Ahmed Moussa, Duaa Shafique, Saleh Ba Raean, and Tarek Hegazy. Construction robotics: Review of intelligent features. *International Journal of Intelligent Robotics and Applications*, 7(3):535–555, 2023. doi: 10.1007/s41315-023-00275-1.

Cindy L Bethel and Robin R Murphy. Review of human studies methods in HRI and recommendations. *International Journal of Social Robotics*, 2 (4):347–359, 2010. doi: 10.1007/s12369-010-0064-9.

Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468. IEEE, 2016.

Thomas Bock. Construction robotics. *Autonomous Robots*, 22:201–209, 2007. doi: 10.1007/s10514-006-9008-5.

Tobias Bonwetsch. *Robotically Assembled Brickwork: Manipulating Assembly Processes of Discrete Elements*. Doctoral Thesis, ETH Zurich, 2015.

C. Borg Costanzi, Z. Y. Ahmed, H. R. Schipper, F. P. Bos, U. Knaack, and R. J. M. Wolfs. 3D Printing Concrete on temporary surfaces: The design and fabrication of a concrete shell structure. *Automation in Construction*, 94:395–404, October 2018. ISSN 0926-5805. doi: 10.1016/j.autcon.2018.06.013.

Cynthia Breazeal. Social interactions in HRI: The robot view. *IEEE transactions on systems, man, and cybernetics, part C (applications and reviews)*, 34(2):181–186, 2004. doi: 10.1109/TSMCC.2004.826268.

Cynthia Brosque, Elena Galbally, Yuxiao Chen, Ruta Joshi, Oussama Khatib, and Martin Fischer. Collaborative welding and joint sealing robots with haptic feedback. In Chen Feng, Thomas Linner, Ioannis Brilakis, Daniel Castro, Po-Han Chen, Yong Cho, Jing Du, Semiha Ergan, Borja Garcia de Soto, Jozef Gašparík, Firas Habbal, Amin Hammad, Kepa Iturralde, Thomas Bock, Soonwook Kwon, Zoubeir Lafhaj, Nan Li, Ci-Jyun Liang, Bharadwaj Mantha, Ming Shan Ng, Daniel Hall, Mi Pan, Wei Pan, Farzad Rahimian, Benny Raphael, Anoop Sattineni, Christian Schlette, Isaac Shabtai, Xuesong Shen, Pingbo Tang, Jochen Teizer, Yelda Turkan, Enrique Valero, and Zhenhua Zhu, editors, *Proceedings of the 38th International Symposium on Automation and Robotics in Construction (ISARC)*, pages 1–8, Dubai, UAE, November 2021. International Association for Automation and Robotics in Construction (IAARC). ISBN 978-952-69524-1-3. doi: 10.22260/ISARC2021/0003.

Reinhard Budde, Karlheinz Kautz, Karin Kuhlenkamp, and Heinz

Züllighoven. What is prototyping? *Information Technology & People*, 6(2/3):89–95, 1990.

Built Robotics. Robots that Build the World. https://www.builtrobotics.com/, 2024.

Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. The museum tour-guide robot RHINO. In *Autonome Mobile Systeme 1998: 14. Fachgespräch Karlsruhe, 30. November–1. Dezember 1998*, pages 245–254. Springer, 1999.

Luis V. Calderita, Araceli Vega, Sergio Barroso-Ramírez, Pablo Bustos, and Pedro Núñez. Designing a cyber-physical system for ambient assisted living: A use-case analysis for social robot navigation in caregiving centers. *Sensors*, 20(14), 2020. ISSN 1424-8220. doi: 10.3390/s20144005.

Luis Vicente Calderita, Luis J Manso, Pablo Bustos, Cristina Suárez-Mejías, Fernando Fernández, and Antonio Bandera. THERAPIST: Towards an autonomous socially interactive robot for motor and neurorehabilitation therapies for children. *JMIR rehabilitation and assistive technologies*, 1 (1):e1, October 2014. ISSN 2369-2529. doi: 10.2196/rehab.3151.

Carlos Campos, Richard Elvira, Juan J Gómez Rodríguez, José MM Montiel, and Juan D Tardós. Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021. doi: 10.1109/TRO.2021.3075644.

Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6015–6022. IEEE, 2019. doi: 10.1109/ICRA.2019.8794134.

Yu Fan Chen, Michael Everett, Miao Liu, and Jonathan P. How. Socially aware motion planning with deep reinforcement learning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1343–1350, September 2017. doi: 10.1109/IROS.2017.8202312.

Yong K Cho, Kinam Kim, Shaojun Ma, and Jun Ueda. A robotic wearable exoskeleton for construction worker's safety and health. In *Construction Research Congress 2018*, pages 19–28, 2018. doi: 10.1061/9780784481288.003.

JO Choi and Dong Bin Kim. A new UAV-based module lifting and transporting method: Advantages and challenges. In *36th International Symposium on Automation and Robotics in Construction (ISARC 2019)*, pages 21–24, 2019.

Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep Reinforcement Learning from Human Preferences. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

Mollie Claypool. Our automated future: a discrete framework for the production of housing. *Architectural Design*, 89(2):46–53, 2019.

Juan Manuel Davila Delgado and Lukumon Oyedele. Robotics in construction: A critical review of the reinforcement learning and imitation learning paradigms. *Advanced Engineering Informatics*, 54:101787, 2022.

Edsger W Dijkstra. A note on two problems in connexion with graphs. In *Edsger Wybe Dijkstra: His Life, Work, and Legacy*, pages 287–290. Springer, 2022.

Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. A survey of embodied ai: From simulators to research tasks. *IEEE Trans-*

actions on Emerging Topics in Computational Intelligence, 6(2):230–244, 2022. doi: 10.1109/TETCI.2022.3141105.

Íñigo Elguea-Aguinaco, Antonio Serrano-Muñoz, Dimitrios Chrysostomou, Ibai Inziarte-Hidalgo, Simon Bøgh, and Nestor Arana-Arexolaleiba. A review on reinforcement learning for contact-rich robotic manipulation tasks. *Robotics and Computer-Integrated Manufacturing*, 81:102517, 2023. doi: 10.1016/j.rcim.2022.102517.

Selen Ercan Jenny, Ena Lloret-Fritschi, Fabio Gramazio, and Matthias Kohler. Crafting plaster through continuous mobile robotic fabrication on-site. *Construction Robotics*, 4:261–271, 2020. doi: 10.1007/s41693-020-00043-8.

Selen Ercan Jenny, Daniela Mitterberger, Ena Lloret-Fritschi, Lauren Vasey, Eliott Sounigo, Ping-Hsun Tsai, Petrus Aejmelaeus-Lindström, David Jenny, Fabio Gramazio, and Matthias Kohler. Robotic on-site adaptive thin-layer printing: Challenges and workflow for design and fabrication of bespoke cementitious plasterwork at full architectural scale. *Architecture, Structures and Construction*, September 2022. ISSN 2730-9894. doi: 10.1007/s44150-022-00062-9.

Emek Erdolu. *Robotic Technologies and the Reconfiguration of Building Construction Work*. Doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA, 2024.

Marina Escobar-Planas, Vicky Charisi, and Emilia Gomez. "That Robot Played with Us!" Children's Perceptions of a Robot after a Child-Robot Group Interaction. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):393:1–393:23, November 2022. doi: 10.1145/3555118.

Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-

based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996.

Michael Everett, Yu Fan Chen, and Jonathan P. How. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, September 2018.

Michael F Everett. *Robot Designed for Socially Acceptable Navigation*. PhD thesis, Massachusetts Institute of Technology, 2017.

Z Fang, Y Wu, A Hassonjee, A Bidgoli, and DC Llach. Towards a distributed, robotically assisted construction framework. In *Proceedings of the 40th Annual Conference of the Association of Computer Aided Design in Architecture (ACADIA)*, pages 320–329, 2020. doi: 10.52842/conf.acadia.2020.1.320.

Benjamin Felbrich, Frederik Wulle, Christoph Allgaier, A Menges, A Verl, KH Wurst, and JH Nebelsick. A novel rapid additive manufacturing concept for architectural composite shell construction inspired by the shell formation in land snails. *Bioinspiration & biomimetics*, 13(2):026010, 2018.

Benjamin; Fru?h Felbrich. Multi-Machine Fabrication: An Integrative Design Process Utilising an Autonomous UAV and Industrial Robots for the Fabrication of Long-Span Composite Structures. In *ACADIA 2017: DISCIPLINES & DISRUPTION [Proceedings of the 37th Annual Conference of the Association for Computer Aided Design in Architecture (ACADIA) ISBN 978-0-692-96506-1] Cambridge, MA 2-4 November, 2017), Pp. 248-259*. CUMINCAD, 2017.

Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window

approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.

Anthony Francis, Claudia Pérez-d'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, et al. Principles and guidelines for evaluating social robot navigation algorithms. *arXiv preprint arXiv:2306.16740*, 2023.

Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596. PMLR, 2018.

A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012.

Sébastien Goessens, Caitlin Mueller, and Pierre Latteur. Feasibility study for drone-based masonry construction of real-scale structures. *Automation in Construction*, 94:458–480, October 2018. ISSN 0926-5805. doi: 10.1016/j.autcon.2018.06.015.

Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improved Techniques for Grid Mapping With Rao-Blackwellized Particle Filters. *IEEE Transactions on Robotics*, 23(1):34–46, February 2007. ISSN 1552-3098. doi: 10.1109/TRO.2006.889486.

David Hahlbrock, Michael Braun, Robin Heidel, Patrik Lemmen, Roland Boumann, Tobias Bruckmann, Dieter Schramm, Volker Helm, and Jan Willmann. Cable Robotic 3D-printing: Additive manufacturing on the construction site. *Construction Robotics*, 6(3):305–318, December 2022. ISSN 2509-8780. doi: 10.1007/s41693-022-00082-3.

Srijeet Halder, Kereshmeh Afsari, Erin Chiou, Rafael Patrick, and Kaveh Akbari Hamed. Construction inspection & monitoring with quadruped robots in future human-robot teaming: A preliminary study. *Journal of Building Engineering*, 65:105814, 2023. doi: 10.1016/j.jobe.2022.105814.

Masashi Hamaya, Takamitsu Matsubara, Tatsuya Teramae, Tomoyuki Noda, and Jun Morimoto. Design of physical user–robot interactions for model identification of soft actuators on exoskeleton robots. *The International Journal of Robotics Research*, 40(1):397–410, 2021. doi: 10.1177/0278364919853618.

Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. doi: 10.1109/TSSC.1968.300136.

Martin Heidegger. *Being and Time*. Suny Press, 2010.

Armin Hornung, Kai M. Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*, 2013. doi: 10.1007/s10514-012-9321-0.

Peiyuan Jiang, Daji Ergu, Fangyao Liu, Ying Cai, and Bo Ma. A Review of Yolo algorithm developments. *Procedia computer science*, 199:1066–1073, 2022. doi: 10.1016/j.procs.2022.01.135.

Ryan Luke Johns, Martin Wermelinger, Ruben Mascaro, Dominic Jud, Ilmar Hurkxkens, Lauren Vasey, Margarita Chli, Fabio Gramazio, Matthias Kohler, and Marco Hutter. A framework for robotic excavation and dry stone construction using on-site materials. *Science Robotics*, 8(84):eabp9758, 2023. doi: 10.1126/scirobotics.abp9758.

Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, et al. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*, 2018.

Shyam Sundar Kannan, Ahreum Lee, and Byung-Cheol Min. External human-machine interface on delivery robots: Expression of navigation intent of the robot. In *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, pages 1305–1312. IEEE, 2021. doi: 10.1109/RO-MAN50785.2021.9515408.

Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7): 846–894, 2011. doi: 10.1177/0278364911406761.

Yusuke Kato, Takayuki Kanda, and Hiroshi Ishiguro. May I help you? Design of human-like polite approaching behavior. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, pages 35–42, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696463.

Kapil Katyal, Yuxiang Gao, Jared Markowitz, Sara Pohland, Corban Rivera, I-Jeng Wang, and Chien-Ming Huang. Learning a group-aware policy for robot navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11328–11335. IEEE, 2022. doi: 10.1109/IROS47612.2022.9981183.

Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620(7976):982–987, 2023. doi: 10.1038/s41586-023-06419-4.

B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2021. doi: 10.1109/TITS.2021.3054625.

Ranjay Krishna, Donsuk Lee, Li Fei-Fei, and Michael S Bernstein. Socially situated artificial intelligence enables learning from human interaction. *Proceedings of the National Academy of Sciences*, 119(39):e2115730119, 2022. doi: 10.1073/pnas.2115730119.

James J Kuffner and Steven M LaValle. RRT-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, volume 2, pages 995–1001. IEEE, 2000. doi: 10.1109/ROBOT.2000.844730.

Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019.

Tessa Lau. Building dusty robotics. *Technology| Architecture+ Design*, 4 (1):14–16, 2020. doi: 10.1080/24751448.2020.1705708.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

Samuel Leder, Ramon Weber, Dylan Wood, Oliver Bucklin, and Achim Menges. Distributed robotic timber construction. *Proceedings of the 39th ACADIA*, 2019. doi: 10.52842/conf.acadia.2019.510.

Samuel Leder, HyunGyu Kim, Metin Sitti, and Achim Menges. Enhanced co-design and evaluation of a collective robotic construction system for

the assembly of large-scale in-plane timber structures. *Automation in Construction*, 162:105390, June 2024. ISSN 0926-5805. doi: 10.1016/j.autcon.2024.105390.

Jiachen Li, Chuanbo Hua, Hengbo Ma, Jinkyoo Park, Victoria Dax, and Mykel J Kochenderfer. Multi-agent dynamic relational reasoning for social robot navigation. *arXiv preprint arXiv:2401.12275*, 2024.

Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2811–2817. IEEE, 2021. doi: 10.1109/ICRA48506.2021.9560769.

Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

Haoyu Liu. *Towards Contextualized Synthetic Dataset for Construction Site*. PhD thesis, Carnegie Mellon University, 2023.

Shuijing Liu, Peixin Chang, Weihang Liang, Neeloy Chakraborty, and Katherine Driggs-Campbell. Decentralized Structural-RNN for Robot Crowd Navigation with Deep Reinforcement Learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3517–3524, Xi'an, China, May 2021. IEEE. ISBN 978-1-72819-077-8. doi: 10.1109/ICRA48506.2021.9561595.

Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In *2023 IEEE International*

Conference on Robotics and Automation (ICRA), pages 2774–2781. IEEE, 2023. doi: 10.1109/ICRA48891.2023.10160968.

R Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. Courier Corporation, 2005.

Wenhan Luo, Junliang Xing, Anton Milan, Xiaoqin Zhang, Wei Liu, and Tae-Kyun Kim. Multiple object tracking: A literature review. *Artificial intelligence*, 293:103448, 2021.

Steve Macenski and Ivona Jambrecic. SLAM Toolbox: SLAM for the dynamic world. *Journal of Open Source Software*, 6(61):2783, 2021. doi: 10.21105/joss.02783.

Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. Core challenges of social robot navigation: A survey, March 2021.

Jack McAlorum, Marcus Perry, Christos Vlachakis, Lorena Biondi, and Brenden Lavoie. Robotic spray coating of self-sensing metakaolin geopolymer for concrete monitoring. *Automation in Construction*, 121:103415, January 2021. ISSN 0926-5805. doi: 10.1016/j.autcon.2020.103415.

Nathan Melenbrink, Justin Werfel, and Achim Menges. On-site autonomous construction robots: Towards unsupervised building. *Automation in Construction*, 119:103312, November 2020. ISSN 0926-5805. doi: 10.1016/j.autcon.2020.103312.

Achim Menges. Fusing the Computational and the Physical: Towards a Novel Material Culture. *Architectural Design*, 85(5):8–15, 2015. ISSN 1554-2769. doi: 10.1002/ad.1947.

Steve De Micoli, Katja Rinderspacher, and Achim Menges. Stone morpholo-

gies: Erosion-based digital fabrication through event-driven control. In *Humanizing Digital Reality*, pages 113–124. Springer, 2018.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv:1312.5602 [cs]*, December 2013. doi: 10.48550/arXiv.1312.5602.

Julio C. Montesdeoca, Marcos Toibero, and Ricardo Carelli. Person-following based on social navigation into the sensorized environments. In *2017 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 799–803, 2017. doi: 10.1109/ROBIO.2017.8324515.

Raul Mur-Artal and Juan D Tardós. Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras. *IEEE transactions on robotics*, 33(5):1255–1262, 2017. doi: 10.1109/TRO.2017.2705103.

Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. ORB-SLAM: A versatile and accurate monocular SLAM system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. doi: 10.1109/TRO.2015.2463671.

R. W. Nuckols, S. Lee, K. Swaminathan, D. Orzel, R. D. Howe, and C. J. Walsh. Individualization of exosuit assistance based on measured muscle dynamics during versatile walking. *Science Robotics*, 6(60):eabj1362, November 2021. doi: 10.1126/scirobotics.abj1362.

NVIDIA. NVIDIA omniverse isaac, 2023.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Patrick Pfaff, Wolfram Burgard, and Dieter Fox. Robust monte-carlo localization using adaptive likelihood models. In *European Robotics Symposium 2006*, pages 181–194. Springer, 2006.

Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.

Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017.

Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. ROS: An open-source robot operating system. In *ICRA Workshop on Open Source Software*, volume 3, page 5. Kobe, Japan, 2009.

Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.

Gilles Retsin. Bits and Pieces: Digital Assemblies: From Craft to Automation. *Architectural Design*, 89(2):38–45, 2019a. ISSN 1554-2769. doi: 10.1002/ad.2410.

Gilles Retsin. Discrete architecture in the age of automation. *Architectural Design*, 89(2), 2019b. doi: 10.1002/ad.2406.

Diego Rodriguez and Sven Behnke. DeepWalk: Omnidirectional bipedal gait by deep reinforcement learning. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3033–3039, 2021. doi: 10.1109/ICRA48506.2021.9561717.

Ahmad EL Sallab, Mohammed Abdou, Etienne Perot, and Senthil Yogamani. Deep reinforcement learning framework for autonomous driving. *arXiv preprint arXiv:1704.02532*, 2017.

Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9339–9347, 2019.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Katie Seaborn, Giulia Barbareschi, and Shruti Chandra. Not only WEIRD but "uncanny"? A systematic review of diversity in Human–Robot Interaction research. *International Journal of Social Robotics*, 15(11):1841–1870, 2023. doi: 10.1007/s12369-023-00968-4.

Tixiao Shan and Brendan Englot. LeGO-LOAM: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4758–4765, 2018. doi: 10.1109/IROS.2018.8594299.

Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5135–5142, October 2020. doi: 10.1109/IROS45743.2020.9341176.

Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. PV-RCNN: Point-Voxel Feature Set Abstraction for 3D Object Detection. In *2020 IEEE/CVF Conference on*

Computer Vision and Pattern Recognition (CVPR), pages 10526–10535, Seattle, WA, USA, June 2020. IEEE. ISBN 978-1-72817-168-5. doi: 10.1109/CVPR42600.2020.01054.

Ben Shneiderman. *Human-Centered AI.* Oxford University Press, 2022.

Sam Shum, Advaith Sethuraman, Zhao Dapeng, and Jean Oh. Social navigation simulator (SNS), February 2022.

David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550(7676):354–359, October 2017. ISSN 1476-4687. doi: 10.1038/nature24270.

Yunlong Song, Mats Steinweg, Elia Kaufmann, and Davide Scaramuzza. Autonomous drone racing with deep reinforcement learning. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1205–1212, September 2021. doi: 10.1109/IROS51168.2021.9636053.

Benjamin Stoler, Meghdeep Jana, Soonmin Hwang, and Jean Oh. T2FPV: Constructing High-Fidelity First-Person View Datasets From Real-World Pedestrian Trajectories, September 2022.

Lucy Suchman. Making work visible. *Communications of the ACM*, 38(9): 56–64, 1995. doi: 10.1145/223248.223263.

Lucy Suchman. Located accountabilities in technology production. *Scandinavian journal of information systems*, 14(2):7, 2002.

Lucy Suchman. Situational awareness: Deadly bioconvergence at the boundaries of bodies and machines. *MediaTropes*, 5(1):1–24, 2015.

Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT press, 2018.

Andreas Thoma, Arash Adel, Matthias Helmreich, Thomas Wehrle, Fabio Gramazio, and Matthias Kohler. Robotic fabrication of bespoke timber frame modules. In *Robotic Fabrication in Architecture, Art and Design 2018: Foreword by Sigrid Brell-Çokcan and Johannes Braumann, Association for Robots in Architecture*, pages 447–458. Springer, 2019. doi: 10.1007/978-3-319-92294-2_34.

Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B Cremers, Frank Dellaert, Dieter Fox, Dirk Hahnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, et al. MINERVA: A second-generation museum tour-guide robot. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No. 99CH36288C)*, volume 3. IEEE, 1999. doi: 10.1109/ROBOT.1999.770401.

Elena Torta, Raymond H. Cuijpers, and James F. Juola. Design of a Parametric Model of Personal Space for Robotic Social Navigation. *International Journal of Social Robotics*, 5(3):357–365, August 2013. ISSN 1875-4805. doi: 10.1007/s12369-013-0188-9.

Peter Trautman and Andreas Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 797–803, October 2010. doi: 10.1109/IROS.2010.5654369.

Nathan Tsoi, Mohamed Hussein, Jeacy Espinoza, Xavier Ruiz, and Marynel Vázquez. Sean: Social environment for autonomous navigation. In *Proceedings of the 8th International Conference on Human-Agent Interaction*, pages 281–283, 2020. doi: 10.1145/3406499.3418760.

Nathan Tsoi, Alec Xiang, Peter Yu, Samuel S. Sohn, Greg Schwartz, Sub-

ashri Ramesh, Mohamed Hussein, Anjali W. Gupta, Mubbasir Kapadia, and Marynel Vázquez. SEAN 2.0: Formalizing and Generating Social Situations for Robot Navigation. *IEEE Robotics and Automation Letters*, 7(4):11047–11054, October 2022. ISSN 2377-3766. doi: 10.1109/LRA.2022.3196783.

Jur van den Berg, Stephen J. Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-Body Collision Avoidance. In Cédric Pradalier, Roland Siegwart, and Gerhard Hirzinger, editors, *Robotics Research*, Springer Tracts in Advanced Robotics, pages 3–19, Berlin, Heidelberg, 2011. Springer. ISBN 978-3-642-19457-3. doi: 10.1007/978-3-642-19457-3_1.

Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7464–7475, 2023.

Ruiqi Wang, Weizheng Wang, and Byung-Cheol Min. Feedback-efficient Active Preference Learning for Socially Aware Robot Navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11336–11343, October 2022. doi: 10.1109/IROS47612.2022.9981616.

Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8:279–292, 1992.

Greg Welch, Gary Bishop, et al. An introduction to the Kalman filter. 1995.

Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256, 1992.

Jan Willmann, Michael Knauss, Tobias Bonwetsch, Anna Aleksandra Apolinarska, Fabio Gramazio, and Matthias Kohler. Robotic timber construc-

tion — Expanding additive fabrication to new dimensions. *Automation in Construction*, 61:16–23, January 2016. ISSN 0926-5805. doi: 10.1016/j.autcon.2015.09.011.

Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017.

Yuning Wu, Jieliang Luo, and Hui Li. Learning dense reward with temporal variant self-supervision. In *ICRA 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, 2022. doi: 10.48550/arXiv.2205.10431.

Yuning Wu, Emek Erdolu, Jiaying Wei, et al. Robot in the loop: An ethnography-based approach to robotics and reinforcement learning to support on-site construction work. *Research Square*, July 8 2024a. URL `https://doi.org/10.21203/rs.3.rs-4509610/v1`. Preprint (Version 1).

Yuning Wu, Jiaying Wei, Jean Oh, and Daniel Cardoso Llach. Towards human-centered construction robotics: A Reinforcement Learning-Driven Companion Robot for Contextually Assisting Carpentry Workers, March 2024b.

Zheng Wu, Wenzhao Lian, Vaibhav Unhelkar, Masayoshi Tomizuka, and Stefan Schaal. Learning dense rewards for contact-rich manipulation tasks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6214–6221. IEEE, 2021. doi: 10.1109/ICRA48506.2021.9561891.

Bo Xiao, Chen Chen, and Xianfei Yin. Recent advancements of robotics in construction. *Automation in Construction*, 144:104591, 2022. doi: 10.1016/j.autcon.2022.104591.

Maria Yablonina, Marshall Prado, Ehsan Baharlou, Tobias Schwinn, and ACHIM Menges. Mobile robotic fabrication system for filament structures. *Fabricate: Rethinking Design and Construction*, 3, 2017.

Xue Yan and PENG Gu. A review of rapid prototyping technologies and systems. *Computer-aided design*, 28(4):307–318, 1996.

LI Yang, ZHOU Huixing, and WANG Shun. Design and Simulation of Ground Leveling Construction Robot. *Journal of Beijing Institute of Civil Engineering & Architecture*, 37(2), 2021.

Ji Zhang and Sanjiv Singh. LOAM: Lidar Odometry and Mapping in Real-time. In *Robotics: Science and Systems X*. Robotics: Science and Systems Foundation, July 2014. ISBN 978-0-9923747-0-9. doi: 10.15607/RSS.2014.X.007.

Ketao Zhang, Pisak Chermprayong, Feng Xiao, Dimos Tzoumanikas, Barrie Dams, Sebastian Kay, Basaran Bahadir Kocer, Alec Burns, Lachlan Orr, Talib Alhinai, et al. Aerial additive manufacturing with multiple autonomous robots. *Nature*, 609(7928):709–717, 2022. doi: 10.1038/s41586-022-04988-4.

Yifu Zhang, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International Journal of Computer Vision*, 129:3069–3087, 2021. doi: 10.1007/s11263-021-01513-4.

Worker Interview Summary

Purpose of System

Interviewer 1: First, let's recap the purpose of the system we designed. We'd like to hear your experience with this machine. You've seen it running, you've seen it working. Especially, you also interacted with the remote controller. So, your interaction with the controller, loading it, unloading it, or delivery, or sending the robot to points, right?

Feedback - Interaction

Worker 2: Sending the robot from point to point was easy. A click of two buttons. Very nice. As long as you have the zones set up, that would be great.

Worker 1: Is that what you guys were doing earlier in the week, setting up all the work points so the remote knew where to send it?

Interviewer 1: Yes, but currently at the prototyping stage, we set it up for the workers.

Feedback - Loading/Unloading

Interviewer 1: How was the loading/unloading experience?

W2: It was all right, a nice bin. Not too congested or clogged, just the right size.

Interviewer 1: The container is temporary. It could be bigger in a future iteration.

W1: Can it be made bigger to transport more materials or equipment?

Interviewer 2: The driving algorithms and sensing are the same. We just need a bigger chassis.

W1 and W3: Approved.

W3: Loading/unloading is relatively simple. You could load a five-gallon bucket the same way or a wheelbarrow. As far as transferring materials from point A to point B, it's relatively simple. Send it on its way.

Additional Use Cases

Interviewer 1: What aspects of this machine do you find successful, and what can be improved? What other functions can be added to a navigating robot on a construction site?

W3:

Manipulator / Tooling: Attachments could enable different functions. Like heavy equipment, which can dig, compact, and drill with different attachments.

Cleaning: A broom attachment could clean dust and dirt, eliminating the need for a person to push a broom. It could work after hours, running remotely, picking up objects, and dumping the load at a specific point.

Layout: An attachment for interior wall layout could triangulate and label walls based on drawings. It could work after hours, labeling and data gathering.

Survey / Layout Before Excavation: Before excavation, the robot could navigate and provide points.

Weaknesses and Improvements

Interviewer 1: What is a weakness in the system? What doesn't really work?

W3: Transferring small pieces of hardware is not the primary concern in this rough phase of construction. Moving heavy items like drywall panels would be helpful for safety and efficiency. The robot could have attachments for different functions.

Interviewer 1: Moving small items is useful for reducing physical burden. For example, in carpentry formwork, workers do more walking for hardware than panels.

W3: If the robot does all the moving, workers can be more productive.

W1: My concern is the robot can't load itself. Is there someone to load it, or would it be quicker to do it manually?

Interviewer 1: We currently don't have an arm, but we are considering it for the future.

Perception and Interaction

Interviewer 2: How would you perceive the image or design of the robot?

Interviewer 1: Did you feel intimidated or find it difficult?

W2: No, it was easy. It looked like a robot.

W3: It's similar to an remote control car.

Interviewer 1: Any other forms of interaction you imagine for this machine?

W1: An app.

W2: An app, but not verbal commands due to the loud environment. An app would work better.

W1: Maybe on a watch.

Interviewer 2: Would you find it more convenient on a watch?

W1 and W2: Yes, that would be great.

Additional Tasks and Safety

Interviewer 1: In which other tasks can the current machine be useful? We saw its potential in formwork panel installation and framing.

W2: Possibly furniture assembly during interior finishing stage when the rough construction is completed. It could work in finished spaces, moving small metal pieces, modular cubical furniture, doors, and hardware, reducing physical strain.

Interviewer 1: Any safety concerns for the robot on-site?

W2: Maybe if it was brighter colored, like a vest, to be seen by equipment operators.

W3: Blinking light.

Interviewer 1: How did you find the speed?

Interviewer 2: We set it to low intentionally.

W2: That makes sense. Different speeds for different times. People will just walk around it in close distance.

Interviewer 2: The robot uses AI to avoid bumping into anyone or causing obstructions.

W3: I observed it doing that around Interviewer 2.

Challenges and Adaptations

Interviewer 1: What can be rearranged on sites to accommodate the robot? What adjustments are realistic?

Interviewer 2: Would you prefer Boston Dynamics legged robots or cleaner floors for wheeled robots?

W2: The first option, but materials will still be on the project.

W3: The building must be kept clean. The robot can scan a barcode and deliver items.

Interviewer 1: You still need some clearance for traffic.

W3: In some buildings, it would work. In this building, it may not. Smaller robots would work better.

Interviewer 2: That is possible.

W3: Payload should be enough, and sensors should work like a truck with a trailer.

Interviewer 1: Any other additions?

W3: Maybe the materials are already delivered on the attachment.

Interviewer 1: Thank you for your time.

End of the conversation.

Related Publications

Journal Articles

- Wu, Yuning*, Erdolu, Emek*, Wei, Jiaying, Oh, Jean, Cardoso Llach, Daniel. **Robot in the Loop: An Ethnography-Based Approach to Robotics and Reinforcement Learning to Support On-Site Construction Work**. Under review. Submitted to *Construction Robotics, special issue on Robotics and AI*, (Springer Nature).

Conference Papers

- Wu, Yuning, Wei, Jiaying, Oh, Jean, Cardoso Llach, Daniel. **Towards Human-Centered Construction Robotics: A Reinforcement Learning-Driven Companion Robot For Contextually Assisting Construction Workers**. Accepted to *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Abu Dhabi, 2024.

- Fang, Zhihao, Wu, Yuning, Hassonjee, Ammar, Bidgoli, Ardavan, Cardoso-Llach, Daniel. **Towards a Distributed, Robotically Assisted Construction Framework: Using Reinforcement Learning to Support Scalable Multi-Drone Construction in Dynamic Environments**. In *Proceedings of the 40th Annual Conference of the Association of Computer Aided Design in Architecture (ACA-*

DIA), vol. 1, pp. 320–329, 2020. doi:10.52842/conf.acadia.2020.1.320.

- Wu, Yuning, Luo, Jieliang, Li, Hui. **Learning Dense Reward with Temporal Variant Self-Supervision**. In *IEEE International Conference on Robotics and Automation (ICRA) 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, Philadelphia, 2022.