

Browsing the Latent Space: A New Approach to Interactive Design Exploration for Volumetric Generative Systems

Tomas Cabezon Pedroso
tcabezon@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Daragh Byrne
daraghb@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA




Figure 1: Left: Latent space Visualization. Center: Interpolation of four chair models. Right: View of the resulting 3D chair.

ABSTRACT

We introduce an interactive tool called "Browsing the Latent Space" (BLS), which explores the potential of latent space visualization and interpolation in 3D generative systems for design space exploration. By visualizing and interpreting the abstract space in which the outputs produced by generative systems operate (the latent space), users gain a better understanding of how the model generates its outputs, and can explore underlying patterns and structures in the data. Using the GET3D generative model, this demonstration showcases how latent space manipulation can enhance the creative process and provides an intuitive and interactive way to explore, remix, and create new chair designs. Allowing interaction with and manipulation of the generative system's outputs in real-time, this tool supports increased awareness of the latent space as a resource for a designer's understanding of generative systems and creative exploration.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
C&C '23, June 19–21, 2023, Virtual Event, USA
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0180-1/23/06.
<https://doi.org/10.1145/3591196.3596815>

CCS CONCEPTS

• Human-centered computing → User interface programming.

KEYWORDS

design exploration, design creativity, latent space, interpolation, interactive visualization, generative systems

ACM Reference Format:

Tomas Cabezon Pedroso and Daragh Byrne. 2023. Browsing the Latent Space: A New Approach to Interactive Design Exploration for Volumetric Generative Systems. In *Creativity and Cognition (C&C '23)*, June 19–21, 2023, Virtual Event, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3591196.3596815>

1 INTRODUCTION

Generative systems have become increasingly popular in recent years, with applications spanning various domains, from art and design to music and literature [2, 4, 7]. However, interactions with these systems are often limited to running the model and obtaining a random and unpredictable result, this is, an output produced by the generative system using a random latent vector without any intentional input from the user. In this demonstration, we inquire about a different approach to interact with generative systems through the interpretation and visualization of the latent space and the use of latent space interpolation as a tool for design exploration

with generative systems giving the user more understanding and control over the outputs.

In this experiment, we focus on the design of chairs as an example of a creative process that has historically relied on remixing previous works [6, 12]. Drawing inspiration from this remix culture within a specific design practice, we introduce a method to enhance the creative process by providing a more intuitive and interactive way to explore, remix, and create new designs using generative volumetric systems. Our approach interprets and visualizes the latent space, revealing the underlying patterns and structures within the data to a designer, while interpolation of this design space enables the generation of meaningful outputs situated between elements within it. Previous work has shown how feature editing is possible within the latent space [8, 10] or how latent space visualization helps reveal significant connections between data points [5] both useful strategies for design exploration.

A latent space refers to a statistical representation of the underlying structure or patterns in the data that the generative system generates. Also known as the feature space, this space is learned during a machine learning training process by mapping the training data to a lower-dimensional vector space. Latent space interpolation, therefore, allows to manipulate the different features of the chairs extracted by the model during training. This not only gives control over the final outputs of the generative system, like color, shape, legs... but also gives the designer an understanding of the way the system interprets what a chair is based on the features that were extracted.

In addition to showcasing the value of latent space visualization and interpolation in generative design, this experiment aims to demonstrate the interaction between users and generative systems through the creation of a working prototype. This is intended to showcase the potential for design exploration of generative systems. For the prototype, we use the recently released GET3D [3] Generative Adversarial Network (GAN) model that Nvidia presented at the last NeurIPS 2022 conference. Generative systems have been gaining popularity in the design community due to their ability to generate high-quality designs in different domains from 2D images to 3D objects like the ones in this demonstration.




Figure 2: Comparison of the traditional Generative Systems' pipeline compared with the pipeline proposed in this demonstration.

2 INTERACTION

There are three main elements of the interactive BLS web tool: the latent space visualization, the interpolation tool, and the resulting 3D chair preview. The three parts of the interaction, as depicted in Figure 2, move from a larger more general visualization of the generative system — that depicts the underlying structure of the data (the latent space, which is a vector space)— to a more specific and focused depiction of an individual model of a chair (corresponding to a latent vector). Technical details of how each part of the tool was developed are offered in the next section.

The interaction starts with a latent space visualization, where the user can interact by zooming and panning through all of the chairs generated with GET3D. A dimensionality reduction algorithm is used to generate a visualizable two-dimensional latent space. This dimensionally reduced space shows how chairs are located and grouped in the latent space based on their texture latent vector. This means that similar texture chairs are located closer to each other.

The user is able to click on the 'Interpolation Tool' button. This opens the interpolation tool. The user can pick 4 chairs from the 5 chair models already pre-prepared for this function¹. Once the models are chosen, a 10x10 grid of chairs is shown with the interpolated elements. The user can hover over this grid to see a larger image of each of the chairs and explore how the features of the chair gradually change from one element of the grid to another. When clicking on any element of the grid, an interactive 3D view of the picked chair is shown, allowing the user to inspect it from any angle.

3 TECHNICAL IMPLEMENTATION

The implementation of the approach involves generating and visualizing the latent space of the GET3D generative system. GET3D is a Generative Adversarial Network (GAN) that consists of two different networks, a generator, that creates fake data samples that aim to be indistinguishable from training samples by inferring the statistical distribution behind the training data, and a discriminator, that evaluates the authenticity of the generated samples. For this work, we only use the generator, as we did not train a model and instead used an already pre-trained model published by Nvidia. We used the generator to map a normal distribution of random vectors — the latent vectors — to their corresponding 3D chairs.

3.1 Generating the Latent Space

To generate the latent space, we need to run the generative system to generate the different chair models. For each chair output by the generative system, we need to provide two random latent vectors that are used by the generator as the input to generate a 3D chair. One input vector for the texture, and another one for the geometry or shape of the chair. Therefore, for each pair of input vectors, we have one output chair (see Figure 3).

3.1.1 Generating the 3D Models, Data Gathering. To generate the latent space, a total of 256 random pair vectors were used to generate the 3D models of chairs that are used in this demonstration. In order

¹Due to memory restrictions of the web application, only five different chair objects are selected for this initial version of the interpolation tool.




Figure 3: Diagram of the generation of a chair in the GET3D ML model.

to make sure that the whole latent space is represented, we used random vectors from a normal distribution to generate the chairs. For the texture of each chair, a random vector of size 9×512 is used, as this is the size of the texture latent space. For the shape, a random vector of size 22×512 is used.




Figure 4: Diagram of the generation of the elements in the latent space for this demonstration.

3.1.2 Generating the latent space, Data visualization. As high-dimensional data is difficult to visualize, and each point in the texture latent space has a size of 9×512 , we used the t-SNE dimensionality reduction algorithm [11] to generate a two-dimensional representation of the latent space. To do this, we flattened each latent vector by reshaping it into a one-dimensional array of size 4,608 and then we use t-SNE to reduce it into two dimensions. For this visualization, we decided to use the texture latent space. Each point in the visualization corresponds to a chair model generated by the generative model, with the texture latent vector used to position the point in the latent space. To provide additional context and meaning to the visualization, we then associated each point with an image of the corresponding chair as shown in the diagram of Figure 5.

Figure 6 is the resulting visualization of the latent space. This visualization shows how, on the left side, the chairs with more metallic or white textures are located. While on the top left, more red textures can be found, and if we move down, more green or yellow ones. This clustering illustrates that the GET3D model is able to discern the relationships between different textures and generate proper outputs accordingly. In fact, in the next part, the interpolation between different texture features can be explored, as well as between different geometry features, showing the understanding of chair features that the generative system learned during training.

3.2 Interpolation

Interpolation is a common mathematical technique that involves finding a point that lies between two given points. In the context




Figure 5: Diagram of the generation and visualization of the latent space into a two-dimensional map.




Figure 6: Resulting two-dimensional representation of the texture latent space.

of generative design, interpolation can be used to manipulate the latent space of a generative system in order to create new designs by combining the features of the points being interpolated [9]. This interpolation generates chair designs that didn't exist in the dataset but instead represent an amalgam of those interpolated.

This interpolation tool allows the exploration of the features learned by the generative system. The intermediate chairs will show the progressive change of the features. Because GET3D has two input latent vectors, one that corresponds to the shape features and the other one that corresponds to the texture features, the resulting interpolated chairs will be interpolating both of these latent vectors and show gradual changes between the shape and texture features of the chairs. Figure 8 shows two different linear interpolations between two chair models.




Figure 7: Diagram of the interpolation process in the latent space between 4 different chair models.



Figure 8: Resulting chairs of the linear interpolation between chair designs in the latent space.

3.3 Model Visualization

When the user clicks on one of the intermediate interpolation outputs, a 3D view of the chair is loaded that lets the user explore the generated output of the generative system, instead of just an image of the chair. To create this visualization, the three.js [1] javascript library is used to load the shape output of the GET3D model (a .obj file) and the texture output of the model (a .png file).

4 CONCLUSION AND FUTURE WORK

This demonstration showcases a web-based tool, "Browsing the Latent Space" (BLS), that enables users to explore the underlying structures and design relationships within generative systems, specifically Generative Adversarial Networks (GANs). We use the example of the GET3D chair generation model to showcase these possibilities. Our experiment shows the potential of latent space visualization and interpolation in generative systems to provide an intuitive and interactive way to explore, remix, and create new chair designs while getting a better understanding of the way generative systems operate and how latent spaces are generated. We believe this approach could be applied as a creativity support tool for designers or as a tool to support design education. Future work will examine these possibilities.

The authors are aware that this demonstration uses pre-trained data. This could be improved in future iterations by using cloud model inference. This change would increase processing time challenging real-time interactions. While our results are promising, our experiment is limited to a single generative system and the a-priori specification of the number of elements of the latent space. Future work could explore how this approach can be applied to other generative systems, as well as developing more sophisticated algorithms for dimensionality reduction that allow bi-directional projections like UMAP or three-dimensional data visualizations.

REFERENCES

- [1] [n. d.]. Three.js – JavaScript 3D Library. <https://threejs.org/>
- [2] Mar Canet Sola and Varvara Guljajeva. 2022. Dream Painter: Exploring creative possibilities of AI-aided speech-to-image synthesis in the interactive art context. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 5, 4 (Sept. 2022), 1–11. <https://doi.org/10.1145/3533386>
- [3] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. 2022. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images. arXiv:2209.11163 [cs.CV]
- [4] Jimpei Hitsuware, Yoshiyuki Ueda, Woojin Yun, and Michio Nomura. 2023. Does human-AI collaboration lead to more creative art? Aesthetic evaluation of human-made and AI-generated haiku poetry. *Computers in Human Behavior* 139 (Feb. 2023), 107502. <https://doi.org/10.1016/j.chb.2022.107502>
- [5] Yang Liu, Eunice Jun, Qisheng Li, and Jeffrey Heer. 2019. Latent Space Cartography: Visual Analysis of Vector Space Embeddings. *Computer Graphics Forum* 38, 3 (June 2019), 67–78. <https://doi.org/10.1111/cgf.13672>
- [6] Eduardo Navas, Owen Gallagher, and Xtine Burrough (Eds.). 2017. *The Routledge companion to remix studies*. Routledge, Taylor and Francis Group, New York ; London.
- [7] Aadars Padiyath and Brian Magerko. 2021. desAIner: Exploring the Use of "Bad" Generative Adversarial Networks in the Ideation Process of Fashion Design. In *Creativity and Cognition*. ACM, Virtual Event Italy, 1–3. <https://doi.org/10.1145/3450741.3466636>
- [8] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. 2020. Interpreting the Latent Space of GANs for Semantic Face Editing. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Seattle, WA, USA, 9240–9249. <https://doi.org/10.1109/CVPR42600.2020.00926>
- [9] Łukasz Struski, Michał Sadowski, Tomasz Danel, Jacek Tabor, and Igor T. Podolak. 2023. Feature-Based Interpolation and Geodesics in the Latent Spaces of Generative Models. *IEEE Transactions on Neural Networks and Learning Systems* (2023), 1–15. <https://doi.org/10.1109/TNNLS.2023.3251848> Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- [10] Meliksah Turker, Alara Dirik, and Pinar Yanardag. 2022. MIDISpace: Finding Linear Directions in Latent Space for Music Generation. In *Creativity and Cognition (C&C '22)*. Association for Computing Machinery, New York, NY, USA, 420–427. <https://doi.org/10.1145/3527927.3532790>
- [11] Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research* 9 (Nov. 2008), 2579–2605.
- [12] D. Wiles and Ch Dymkowski. 2013. *1000 Chairs*. Cambridge University Press.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009