

Making Physical Objects with Generative AI and Robotic Assembly: Considering Fabrication Constraints, Sustainability, Time, Functionality and Accessibility

ALEXANDER HTET KYAW, Massachusetts Institute of Technology, United States

SE HWAN JEON, Massachusetts Institute of Technology, United States

MIANA SMITH, Massachusetts Institute of Technology, United States

NEIL GERSHENFELD, Massachusetts Institute of Technology, United States

3D generative AI enables rapid and accessible creation of 3D models from text or image inputs. However, translating these outputs into physical objects remains a challenge due to the constraints in the physical world. Recent studies have focused on improving the capabilities of 3D generative AI to produce fabricable outputs, with 3D printing as the main fabrication method. However, this workshop paper calls for a broader perspective by considering how fabrication methods align with the capabilities of 3D generative AI. As a case study, we present a novel system using discrete robotic assembly and 3D generative AI to make physical objects. Through this work, we identified five key aspects to consider in a physical making process based on the capabilities of 3D generative AI. 1) Fabrication Constraints: Current text-to-3D models can generate a wide range of 3D designs, requiring fabrication methods that can adapt to the variability of generative AI outputs. 2) Time: While generative AI can generate 3D models in seconds, fabricating physical objects can take hours or even days. Faster production could enable a closer iterative design loop between humans and AI in the making process. 3) Sustainability: Although text-to-3D models can generate thousands of models in the digital world, extending this capability to the real world would be resource-intensive, unsustainable and irresponsible. 4) Functionality: Unlike digital outputs from 3D generative AI models, the fabrication method plays a crucial role in the usability of physical objects. 5) Accessibility: While generative AI simplifies 3D model creation, the need for fabrication equipment can limit participation, making AI-assisted creation less inclusive. Discrete robotic assembly with 3D generative AI shows promising results in addressing key aspects such as fabrication constraints, time, sustainability, and functionality, though accessibility remains a limitation. These five key aspects provide a framework for assessing how well a physical making process aligns with the capabilities of 3D generative AI and values in the world.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**.

Additional Key Words and Phrases: Generative AI, Robotic Assembly, Digital Fabrication, Sustainability, Time, Functionality

1 Introduction

Recent advances in 3D generative AI have enabled the fast and easy creation of digital models from text descriptions or image inputs [10]. For example, models like InstantMesh, Latte3D, GET3D, and Shap-E allow users to generate 3D models in seconds [27] [26] [8] [12]. However, while AI-generated 3D models are easy to produce, fabricating them into physical objects remains a challenge due to real-world constraints [6] [16]. Previous research has focused on improving capabilities of generative AI to produce fabricable outputs, with 3D printing as the primary fabrication method [5] [3] [7] [17]. However, 3D printing as a fabrication method has significant limitations [28]. While it allows for precise manufacturing, it can take hours or even days to produce a single object, making real-time human-AI collaboration more challenging [13] [29] [4]. Additionally, desktop 3D printing is best suited for small to medium-sized and tabletop-scale objects, often constrained by the size of the printer bed [2] [21]. This limitation prevents the fabrication of larger structures, such as furniture or room-scale objects. As generative AI continues to advance, fabrication methods should align with its speed and scale, enabling on-demand, large-scale production in a sustainable manner. Relying on 3D printing or only improving the capabilities of 3D generative AI models is insufficient to fully realize the potential of

Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). Copyright remains with the author(s).




Fig. 1. Figure 1. Demonstration of 3D Generative AI to Robotic Assembly from a User Prompt.




Fig. 2. System pipeline from speech to robotic assembly

generative AI in physical making. This paper calls for a broader perspective by considering how the making process aligns with the capabilities of generative AI.

To address this, we explore discrete robotic assembly as an alternative fabrication process for 3D generative AI outputs. By leveraging generative AI, our approach lowers the barrier to entry for robotic assembly, making it possible for users to speak objects into reality without programming or 3D modeling [1]. Through this case study, we propose a framework for evaluating fabrication methods based on five key aspects: fabrication constraints, production time, sustainability, functionality, and accessibility. The workshop paper contributes to bridging the gap between generative AI and physical making by examining how fabrication processes can better support AI-generated designs. We propose a novel system using 3D generative AI and discrete robotic assembly as an on-demand, prompt-to-production method with the goal of creating functional objects in a sustainable and accessible manner. (Fig 1)

2 Method

In order to convert speech into physical objects, we present a system that uses discrete robotic assembly and 3D generative AI, employing a 6-axis robotic arm to assemble smaller modular components into larger objects. In our system, we use the UR10 robot from Universal Robots and Meshy.AI as the text-to-3D generative AI model [24][18] (Fig 2). Each paragraph investigates 3D generative AI-based discrete robotic assembly through five key aspects: fabrication constraints, production time, sustainability, functionality, and accessibility.

Fabrication Constraints: 3D generative AI models typically output meshes or point clouds [15] [20]. However, they do not inherently consider fabrication constraints or generate component-level representations suitable for robotic assembly. Additionally, AI models can generate a wide variety of 3D geometries, requiring the making process to adapt to the variability of generative AI outputs. To address this, we discretize the AI-generated geometry into assembly components through a voxelization algorithm based on the dimensions of the assembly component. While this provides a component-level representation for robotic assembly, it does not guarantee that the robot will successfully complete the

assembly. To ensure assembly feasibility, we further modify the geometry based on algorithmic checks and fabrication constraints, including overhang detection, connectivity search, and robotic arm reachability analysis.

Time: While generative AI can produce 3D models in seconds, manufacturing physical objects takes considerably longer. Faster assembly could enable a closer iterative design loop between humans and AI in the making process in the future [9]. Unlike traditional manufacturing methods, discrete robotic assembly relies on prefabricated modular parts, reducing the need for material preprocessing. However, we acknowledge that other methods, such as 3D printing, also require material preparation (e.g., filament spooling, resin curing). For a fair comparison, we measure fabrication time from the user’s request, excluding material preparation and prefabrication.

Sustainability: 3D Generative AI enables the creation of countless assets, but translating this capability into the physical world can be resource-intensive, unsustainable, and irresponsible [25]. We study how discrete robotic assembly can mitigate material waste by enabling the disassembly and reassembly of components into different objects. We start with 40 modular components and evaluate how many distinct objects can be assembled from the same set of parts. The goal of this method is to promote sustainable fabrication by ensuring that physical artifacts generated through AI can be efficiently repurposed rather than discarded after a single use.

Functionality: As text-to-3D models advance, it is also crucial from the fabrication process to be able to make objects that are functional. For example, if the AI generates a chair, the robotic system may successfully assemble a chair-like structure. However, if this object lacks structural integrity and collapses when someone sits on it, it could pose a safety risk. In our case study, we assemble lattice-based cuboidal components with magnets on each face for secure yet reversible connections [11]. Our goal is to create functional objects that can withstand some amount of compression.

Accessibility: While generative AI offers an accessible method for creating 3D models, physical fabrication often requires specialized equipment and hardware. This system can be implemented using an industrial robotic arm, a low-cost robotic arm, or DIY robotic arms assembled at home from 3D-printed parts. To assess accessibility, we calculate the cost of the industrial robotic arm and alternatives. Expanding accessibility through lower-affordable options can make AI-driven fabrication more inclusive, enabling broader participation [14].

3 Results

Our results show that discretizing AI-generated meshes into component-level representations can handle the variability in generative AI outputs for discrete robotic assembly. (Fig 3) However, simply voxelizing an AI-generated mesh into a component-level representation does not always guarantee successful robotic assembly. Our experiments showed that most attempts failed without overhang detection, connectivity analysis, and robotic arm reachability assessment.

For assembly time, we measured the total duration required to assemble each object, starting from the moment the user prompts the system. The assembled objects have an average volume of around 7500 cm³, with nearly all of them being assembled in under five minutes. Considering that generative AI can produce objects of any size in seconds, we believe that the ability of discrete robotic assembly to construct large objects within minutes presents a promising match. With a fast prompt-to-physical workflow, future research could explore an iterative design loop between humans and generative AI in the making process.

Our results show that discrete robotic assembly enables sustainable production through the reuse of modular components to construct different objects. Using a set of 40 modular components, we successfully assembled seven distinct objects from the same set. This finding highlights the potential for material-conscious and responsible physical making with generative AI. Unlike traditional manufacturing methods, where material waste accumulates after each fabrication cycle, discrete robotic assembly enables disassembly and repurposing.




Fig. 3. A variety of objects created from user prompts using 3D generative AI and discrete robotic assembly

To assess functionality, we check both the shape and usability of the object. We use GPT-4's vision-language capabilities to verify whether the assembled object corresponds to the user prompt [19]. Our experiments with the vision language model indicate that all assembled objects resemble the user prompt. As for the usability of the objects, the chair and stool we assembled could not support sitting, whereas the shelf and table functioned as intended, allowing us to place objects on them. The Letter T and the tall dog were evaluated only based on their visual resemblance using GPT-4.

Our system enables users to create physical objects from speech without requiring expertise in 3D modeling or robotic programming. However, cost is the main accessibility barrier of this system. An industrial UR10 robotic arm used in this study can cost between \$40,000 and \$50,000 [22]. DIY robotic arms constructed from 3D-printed parts or cheaper off-the-shelf robotic arms can cost under \$1,000 but have reduced precision and payload capacity [23] [11]. However, further tests with low-cost or 3D-printed robotic arms are needed, as this study was conducted with a UR10 robotic arm.

4 Conclusion

While current text-to-3D outputs cannot be directly assembled, we demonstrate that discretizing them into component-level representations can accommodate variability in AI-generated designs and enable robotic assembly. We show that overhang detection, connectivity search, and robotic arm reachability analysis are necessary since 3D generative AI outputs do not inherently consider physical constraints. The rapid creation of physical objects aligns with AI's fast generative capabilities, opening new research avenues for future studies on real-time, human-AI iterative design workflows in physical fabrication. Furthermore, reusing components enables a more sustainable and responsible approach to making with generative AI outputs compared to wasteful manufacturing methods. By leveraging generative AI, our approach contributes to lowering the barrier to robotic assembly, enabling users to speak objects into reality. However, challenges remain, particularly in ensuring the functionality of certain objects and accessibility. Future work could enhance the structural integrity and functionality of objects while exploring the integration of generative AI with

discrete robotic assembly across varying component types, and connection methods. This workshop paper contributes to a broader framework on how fabrication methods align with the capabilities of 3D generative AI. Additionally, the study contributes to the advancement of AI-driven fabrication by integrating 3D generative AI with discrete robotic assembly, enabling on-demand, sustainable, and accessible production.

References

- [1] Rafael Ballagas, Jishang Wei, Mithra Vankipuram, Zhiyuan Li, Keanu Spies, and Hiroshi Horii. 2019. Exploring Pervasive Making Using Generative Modeling and Speech Input. *IEEE Pervasive Computing* 18, 4 (Oct. 2019), 20–28. doi:10.1109/MPRV.2019.2929130 Conference Name: IEEE Pervasive Computing.
- [2] Xuelin Chen, Hao Zhang, Jinjie Lin, Ruizhen Hu, Lin Lu, Qixing Huang, Bedrich Benes, Daniel Cohen-Or, and Baoquan Chen. 2015. Dapper: decompose-and-pack for 3D printing. *ACM Trans. Graph.* 34, 6 (Nov. 2015), 213:1–213:12. doi:10.1145/2816795.2818087
- [3] Valdemar Danry and Cenk Güzelis. 2023. Organs without Bodies. Generating physical objects with AI. <https://organs.media.mit.edu/>
- [4] Mustafa Doga Dogan, Patrick Baudisch, Hrvoje Benko, Michael Nebeling, Huaishu Peng, Valkyrie Savage, and Stefanie Mueller. 2022. Fabricate It or Render It? Digital Fabrication vs. Virtual Reality for Creating Objects Instantly. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems (CHI EA '22)*. Association for Computing Machinery, New York, NY, USA, 1–4. doi:10.1145/3491101.3516510
- [5] Kristen M. Edwards, Brandon Man, and Faez Ahmed. 2024. Sketch2Prototype: rapid conceptual design exploration and prototyping with generative AI. *Proceedings of the Design Society* 4 (May 2024), 1989–1998. doi:10.1017/pds.2024.201
- [6] Faraz Faruqi, Ahmed Katary, Tarik Hasic, Amira Abdel-Rahman, Nayeemur Rahman, Leandra Tejedor, Mackenzie Leake, Megan Hofmann, and Stefanie Mueller. 2023. Style2Fab: Functionality-Aware Segmentation for Fabricating Personalized 3D Models with Generative AI. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3586183.3606723
- [7] Faraz Faruqi, Yingtao Tian, Vrushank Phadnis, Varun Jampani, and Stefanie Mueller. 2024. Shaping Realities: Enhancing 3D Generative AI with Fabrication Constraints. In *CHI 2024 Workshop on Generative AI and HCI* arXiv. <http://arxiv.org/abs/2404.10142> arXiv:2404.10142 [cs].
- [8] 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. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=GAUwreODU5L>
- [9] Frederic Gmeiner, Humphrey Yang, Lining Yao, Kenneth Holstein, and Nikolas Martelaro. 2023. Exploring Challenges and Opportunities to Support Designers in Learning to Co-create with AI-based Manufacturing Design Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–20. doi:10.1145/3544548.3580999
- [10] Roberto Gozalo-Brizuela and Eduardo C. Garrido-Merchán. 2023. A survey of Generative AI Applications. <http://arxiv.org/abs/2306.02781> arXiv:2306.02781 [cs].
- [11] Benjamin Jenett, Amira Abdel-Rahman, Kenneth Cheung, and Neil Gershenfeld. 2019. Material-Robot System for Assembly of Discrete Cellular Structures. *IEEE Robotics and Automation Letters* 4, 4 (Oct. 2019), 4019–4026. doi:10.1109/LRA.2019.2930486 Conference Name: IEEE Robotics and Automation Letters.
- [12] Heewoo Jun and Alex Nichol. 2023. Shap-E: Generating Conditional 3D Implicit Functions. <http://arxiv.org/abs/2305.02463> arXiv:2305.02463 [cs].
- [13] Jeeun Kim, Anhong Guo, Tom Yeh, Scott E. Hudson, and Jennifer Mankoff. 2017. Understanding Uncertainty in Measurement and Accommodating its Impact in 3D Modeling and Printing. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*. Association for Computing Machinery, New York, NY, USA, 1067–1078. doi:10.1145/3064663.3064690
- [14] Stacey Kuznetsov and Eric Paulos. 2010. Rise of the expert amateur: DIY projects, communities, and cultures. In *Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries (NordiCHI '10)*. Association for Computing Machinery, New York, NY, USA, 295–304. doi:10.1145/1868914.1868950
- [15] Chenghao Li, Chaoning Zhang, Joseph Cho, Atish Waghvase, Lik-Hang Lee, Francois Rameau, Yang Yang, Sung-Ho Bae, and Choong Seon Hong. 2024. Generative AI meets 3D: A Survey on Text-to-3D in AIGC Era. doi:10.48550/arXiv.2305.06131 arXiv:2305.06131 [cs].
- [16] Liane Makatura, Michael Foshey, Bohan Wang, Felix HahnLein, Pingchuan Ma, Bolei Deng, Megan Tjandrasuwita, Andrew Spielberg, Crystal Elaine Owens, Peter Yichen Chen, Allan Zhao, Amy Zhu, Wil J. Norton, Edward Gu, Joshua Jacob, Yifei Li, Adriana Schulz, and Wojciech Matusik. 2023. How Can Large Language Models Help Humans in Design and Manufacturing? doi:10.48550/arXiv.2307.14377 arXiv:2307.14377 [cs].
- [17] Ryan McClelland. 2022. Generative design and digital manufacturing: using AI and robots to build lightweight instrument structures. In *Current Developments in Lens Design and Optical Engineering XXIII*, Vol. 12217. SPIE, 141–148. doi:10.1117/12.2646476
- [18] Meshy. 2025. Meshy - Free AI 3D Model Generator. <https://www.meshy.ai/>
- [19] OpenAI. 2025. OpenAI Platform. <https://platform.openai.com>
- [20] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2022. DreamFusion: Text-to-3D using 2D Diffusion. In *The Eleventh International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=FjNys5c7VyY>
- [21] Prusa. [n.d.]. Original Prusa MK4S 3D Printer | Original Prusa 3D printers directly from Josef Prusa. <https://www.prusa3d.com/product/original-prusa-mk4s-3d-printer-5/>

- [22] QVIRO. 2025. UR10e Cobot Reviews, Price, Use-cases, Find your cobot on Qviro.com. <https://qviro.com/product/universal-robots/ur10e>
- [23] Elephant Robotics. 2025. Elephant Robotics myCobot 280 for Arduino 6 DOF Collaborative Robot (for Arduino Version). <https://shop.elephantrobotics.com/products/mycobot-280-for-arduino-6-dof-collaborative-robot>
- [24] Universal Robots. 2025. UR10e Medium-sized, versatile cobot. <https://www.universal-robots.com/products/ur10-robot/>
- [25] Valkyrie Savage, Sarah Homewood, and Irina Shklovski. 2023. Objectify: Better Living Through Anticipatory, Just-for-you 3D Printing!. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–8. doi:10.1145/3544549.3582748
- [26] Kevin Xie, Jonathan Lorraine, Tianshi Cao, Jun Gao, James Lucas, Antonio Torralba, Sanja Fidler, and Xiaohui Zeng. 2024. LATTE3D: Large-scale Amortized Text-To-Enhanced3D Synthesis. In *Computer Vision – ECCV 2024: 18th European Conference, Milan, Italy, September 29–October 4, 2024, Proceedings, Part LXXVII*. Springer-Verlag, Berlin, Heidelberg, 305–322. doi:10.1007/978-3-031-72980-5_18
- [27] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. 2024. InstantMesh: Efficient 3D Mesh Generation from a Single Image with Sparse-view Large Reconstruction Models. doi:10.48550/arXiv.2404.07191 arXiv:2404.07191 [cs].
- [28] Nur Yildirim, James McCann, and John Zimmerman. 2020. Digital Fabrication Tools at Work: Probing Professionals’ Current Needs and Desired Futures. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376621
- [29] Jiayi Zhou, Renzhong Li, Junxiu Tang, Tan Tang, Haotian Li, Weiwei Cui, and Yingcai Wu. 2024. Understanding Nonlinear Collaboration between Human and AI Agents: A Co-design Framework for Creative Design. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–16. doi:10.1145/3613904.3642812