

Sketch2Prototype: rapid conceptual design exploration and prototyping with generative AI

Kristen M. Edwards✉, Brandon Man and Faez Ahmed

Massachusetts Institute of Technology, United States of America

✉ kme@mit.edu

Abstract

Sketch2Prototype is an AI-based framework that transforms a hand-drawn sketch into a diverse set of 2D images and 3D prototypes through sketch-to-text, text-to-image, and image-to-3D stages. This framework, shown across various sketches, rapidly generates text, image, and 3D modalities for enhanced early-stage design exploration. We show that using text as an intermediate modality outperforms direct sketch-to-3D baselines for generating diverse and manufacturable 3D models. We find limitations in current image-to-3D techniques, while noting the value of the text modality for user-feedback.

Keywords: prototyping, conceptual design, artificial intelligence (AI), sketching

1. Introduction

During product design and development, a design concept moves through many modalities. It may be represented as a sketch, a textual description, a looks-like or works-like prototype, and, ultimately, the final product (Ulrich et al., 2020). In early-stages of engineering design, sketches and prototypes are common modalities for communicating ideas and exploring the design space (Bao et al., 2018). Developing looks-like prototypes is an essential phase in the engineering design process, offering a tangible, visual representation of a product idea, and communicating the design concept to stakeholders. However, prototyping can be time-consuming and resource-intensive, involving multiple iterations and manual adjustments to achieve the desired outcome (Lauff et al., 2019).

During conceptual design, sketching and prototyping are often done sequentially, but research has shown benefits to performing them in parallel (Bao et al., 2018). Given that conceptual design determines up to 70-80% of a product's lifetime cost (Corbett and Crookall, 1986; Pahl et al., 2007) exploring the design space with proper breadth and depth is a worthwhile cause. However, sketching and prototyping are often done in sequence because sketching is quicker and has lower overhead (Ulrich et al., 2020).



Figure 1. Our Sketch2Prototype framework which takes in a conceptual design sketch as input, and produces multiple inspired images, a 3D model, and finally a fabricated prototype

Recent breakthroughs in generative AI have enabled people to generate novel, unseen images by learning underlying patterns in training data. Through generative models, the looks-like prototyping process can be streamlined, allowing for rapid generation and iteration of design options, thus significantly reducing time and costs associated with manual methods. This enables design space exploration, motivating designers with a diverse set of examples. Moreover, incorporating machine learning in the prototype development process opens the door for enhanced user interaction and feedback through easy iterations, as shown later in Figure 5B. Furthermore, communication between the development team can be bolstered by having a physical manifestation of their design vision. Ultimately, the integration of machine learning in creating looks-like prototypes facilitates a more efficient, cost-effective, and user-centered approach to product development, aligning technological innovation with aesthetic and practical design needs.

In this work, we propose Sketch2Prototype, a framework for understanding sketches, generating new conceptual images inspired by those sketches, converting the images to 3D models, and finally fabricating a looks-like prototype from these 3D models. There are various existing models that can perform each of these subtasks, we demonstrate several state-of-the-art methods. We found that GPT-4V(ision), a vision language model by [OpenAI, 2023](#), is able to interpret and explain hand-drawn sketches ([Picard et al., 2023](#)). Therefore, we use GPT-4V, to convert sketches into textual prompt descriptions, then DALL-E 3, which is a generative text-to-image model, generates a set of more descriptive images from the text. We then use those images to generate a 3D model which, after postprocessing, we fabricated via additive manufacturing. We demonstrate Sketch2Prototype in a series of case studies with real hand drawn sketches of milk frothers, phone stands, pen and coin holders, and mugs. Our method enhances design space exploration by three means: 1) inherent design expansion caused from automatically generating multiple 2D images inspired by one sketch 2) increased breadth and depth of exploration made possible by working with sketches and prototypes in parallel ([Bao et al., 2018](#)), and 3) allowing for user-centered feedback via the text modality. The contributions of this work are as follows:

1. We introduced a generative AI-based framework to rapidly create a prototype from a sketch, enabling design exploration as the design moves through sketch, text, image, and 3D modalities.
2. We compared our framework to direct sketch-to-3D and ControlNet-generated image-to-3D frameworks. Our model generates more diverse images and manufacturable designs.
3. We demonstrated examples of the successful Sketch2Prototype, moving from a hand drawn sketch to a fabricated 3D looks-like prototype for four design categories and six designs.
4. We built an open-source dataset of 1,087 milk frother sketches each with four paired images inspired by the sketch and generated by our framework. This results in 4,348 images.

2. Related works

In the following sections we discuss related works regarding sketching and prototyping in engineering design, recent advancements in vision language models, image-to-3D models, and 3D representations.

2.1. Sketching and prototyping in engineering design

Sketching is documented as a valuable skill in engineering design, and researchers have studied ways to encourage and understand sketching in engineering education ([Schmidt et al., 2012](#); [Das and Yang, 2022](#)). Sketching provides a rapid external representation that comes at very little cognitive cost ([Goldschmidt, 2014](#)). Researchers also explored creativity and decision making with sketches ([Toh and Miller, 2016](#)), using sketching for finite element analysis ([Murugappan et al., 2017](#)), and using machine learning to predict creativity-ratings from sketches and text ([Edwards et al., 2022](#); [Song et al., 2023](#)).

In product design, prototypes can be defined as "an approximation of the product along one or more dimensions of interest" ([Ulrich et al., 2020](#)). Prototypes, as well as the process of building and testing them, offer invaluable information to designers ([Das and Yang, 2022](#)). As such, many works have surveyed and explored different prototyping strategies ([Camburn et al., 2017](#); [Hansen and Ozkil, 2020](#)). Research has explored how combining sketching and prototyping during conceptual design impacts design space exploration. On average, only sketching leads to a broader design space and generated more novel designs, only prototyping leads to more aesthetically pleasing designs with better

functionally, both sketching and prototyping explored and generated final ideas that were perceived as more creative (Bao et al., 2018). Sketching can lead to a higher quantity of designs (Neely et al., 2013), but prototyping can be used to both explore and refine designs (Elverum et al., 2016). Furthermore, this work suggests benefits of using both sketch and prototype modalities during conceptual design to explore the design space with both breadth and depth and ultimately generate creative designs.

One challenge that prevents designers from prototyping in parallel with sketching is that prototypes are often slower to create and have higher associated costs than sketching (Lauff et al., 2019), and designers are often reluctant to spend money and time on things when uncertainty is high, like early in the conceptual stage (Ulrich et al., 2020). This is where we believe the integration of machine learning to efficiently create looks-like prototypes presents a transformative opportunity.

2.2. Large language models and vision language models

Large Language Models (LLMs) are billion-parameter transformers that are pre-trained on significant amounts of data, enabling them to perform a wide variety of natural language processing tasks such as translation, summarization and recognition. LLMs such as LLaMA (Touvron et al., 2023) learn to generate text aligned with human preferences through Reinforcement Learning with Human Feedback. Vision Language Models (VLMs) have also become quite popular. To create cohesive understanding between image and text, CLIP (Contrastive Image-Language Pretraining) (Radford et al., 2023) is a model that creates a joint embedding space between vision and language. CLIP is an efficient metric for measuring similarity between an image-text or image-image pair (Hessel et al., 2022). Hence, VLMs such as GPT-4V and FLAVA (Singh et al., 2022) are similar to LLMs except that they train on multimodal datasets and often employ CLIP learn a joint space between language and vision. VLMs need to understand complex relationships between text and image, thus requiring more data and are also harder to train. While there is some overlap in applications, VLMs are uniquely suited for tasks that require understanding the visual aspects of content, whereas large LLMs excel in purely linguistic tasks. As early-stage design is a highly multimodal process, VLMs are better suited for our applications.

Text-to-image synthesis has been explored via models like Imagen (Saharia et al., 2023), DALL-E 3 (Betker et al., 2023). Many researchers have used diffusion models for text-to-image tasks as diffusion models tend to be faster. Many text-to-image models also enable users control over their picture by prompting the model to change specific regions of an image via text, known as inpainting. However, inpainting for these models is often limited to text.

To enable users to control image generation with their own sketches, models such as ControlNet and T2IAdapter have emerged. These models freeze the text-to-image model and use the users' sketches to guide the text-to-image generation process. However, these models give too much control to the users, and often results in images too similar to the users' sketch, meaning less exploration of the design space, as we show later, these models lead to fixation on the original concept. By leveraging VLMs' ability to provide text descriptions, which can then be expanded by text-to-image models, we generate a highly diverse set of prototypes.

2.3. 3D representations and Image-to-3D models

3D representations enable users to visualize the dimensions and proportions of objects. Furthermore, 3D representations enable manufacturing decisions, by displaying different geometric constraints. Lastly, 3D prototypes enable useful feedback from designers, customers, and stakeholders regarding the user-experience with a design (Das and Yang, 2022). Neural Radiance Fields (NeRFs) (Mildenhall et al., 2020) have become a popular method for 3D scene representations. Although NeRFs are used in 3D reconstruction (Barron et al., 2022) and generation (Poole et al., 2022), optimizing NeRFs are time consuming to train and memory intensive. 3D Gaussian splatting (Kerbl et al., 2023) is a recent alternative to NeRFs and has demonstrated promising results in both speed and quality in 3D reconstruction. Recent work has tried applying Gaussian Splatting to generation tasks (Tang et al., 2023) that outperform methods that use NeRFs for 3D representations.

Image-to-3D models try to generate 3D assets from a single image, which can also be reformulated as a single-view 3D reconstruction tasks, but often produce blurry results (Yu et al., 2021). Using image captioning models, text-to-3D methods can be adapted for image-to-3D generation (Melas-Kyriazi et al.,

2023). Dream Gaussian (Tang et al., 2023) is a recent model that uses companioned mesh extraction and texture refinement in UV space. Even though the results of Dream Gaussian are promising, it fails to produce high quality models of unseen models. Shap-E (Jun and Nichol, 2023) is another recent image-to-3D and text-to-3D model that utilizes a conditional diffusion model to output high fidelity 3D objects. In engineering design, it is often time consuming for designers to create high-quality images that can be used for image-to-3D generation tasks. Sketches are often abstract, lack detail, and are often unfit for image generation. Past work has tried to predict 3D functionality from a 2D image; however, 3D information performs best (Edwards et al., 2021). Our work shows an end-to-end system that generates multiple high-quality images from the original sketch, which can then generate a printable 3D prototype.

3. Methodology

3.1. Framework from sketch to prototype

Our proposed framework treats the Sketch2Prototype problem as a sequence of tasks that move between design modalities: from sketch-to-text, then from text-to-image(s), and finally from image-to-3D model, as shown in Figure 2. For sketch-to-text, we fed our sketch as input along with a verbal description of the sketch into GPT-4V and prompted it to give it a description of the image. We also asked GPT-4V to describe it such that it will be passed as a prompt into DALL-E 3. DALL-E 3 performs text-to-image by converting the text description of the original sketch into a text embedding, then feeding it into a diffusion prior to generate an image embedding, which finally gets decoded into an image. We chose not to extract the words found on the sketches as they may give semantic meaning to specific areas of the image. To generate a variety of novel designs, we ask DALL-E 3 to generate 4 images from the original prompt. DALL-E 3 also attempts to generate more diverse images by rephrasing the input prompt. The resulting image generated from DALL-E 3 may contain text, which negatively affects the generation quality when converting from image to 3D. To prevent this, we manually select a set of images that do not include any text and feed it into our image-to-3D model. Current state-of-the-art models such as Shap-E (Jun and Nichol, 2023) and DreamGaussian have varied performance depending on the provided image. As a result, we feed our images into three state-of-the-art models (One-2-3-45, DreamGaussian and Shap-E) and pick the mesh that is most similar to the original image while also being the most manufacturable.

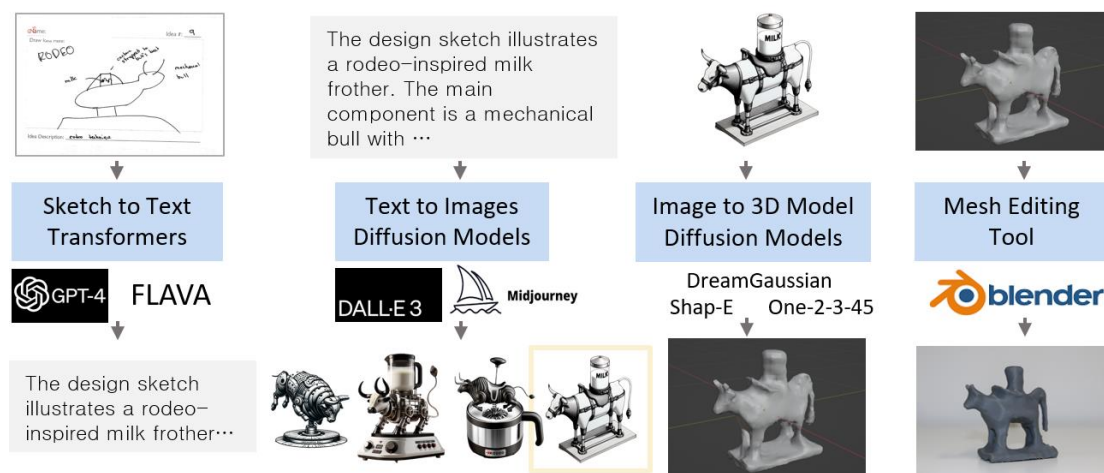


Figure 2. The Sketch2Prototype framework uses transformer-based models for the sketch-to-text and text-to-image steps, as well as an encoder and conditional diffusion model for image-to-3D model. Post-processing of the 3D model is performed in Blender

3.2. Post-processing and fabrication

While models like Shap-E excel at generating 3D models, they may not adhere to fabrication requirements. Hence, we perform post-processing of the 3D model in Blender 3.6.5. Shap-E outputs a 3D model in the Polygon File Format (PLY) family. This can be directly imported to Blender 3.6.5, post-processed as needed, and exported as an STL file, which can be 3D printed. We print the models

with either the Bambu Lab X1-Carbon Combo 3D Printer, which is a fused deposition modelling printer, or Formlabs Form 3 printer, which is a stereolithography printer.

4. Results

In this section, we display the results of our framework via a number of examples. Figure 3 demonstrates how the Sketch2Prototype framework enables exploration of the design space. From one design sketch of a phone stand, our framework leads to three diverse prototypes. The images and 3D printed models show a diverse set of functional phone stands. Figure 4 showcases the full Sketch2Prototype framework for four different design types: a rodeo-inspired milk frother, a pen-and-coin holder, a phone stand, and a backpack-inspired mug. For each of these, we demonstrate an automatic exploration of the design space in the text-to-image step. Here, using a text description created from a sketch via generative-AI, a designer is automatically presented with any number of detailed design images inspired by the sketch. We chose to display four images for each sketch, however there is no imposed limit on this. A benefit of design exploration in this stage is that it mitigates design fixation, and can thus be used as an assistant for designers. Furthermore, showing a designer multiple diverse examples can aid in creative ideation. To assess the diversity and feasibility of our designs, we compare the diversity and manufacturability of our model's designs to those made from a sketch alone or using ControlNet, which adheres to the sketch geometry (Figure 5A).




Figure 3. Our framework enables exploration of the design space by automatically generating multiple diverse images inspired by one sketch. Here a single sketch results in three fabricated looks-like prototypes

4.1. Enhanced diversity and manufacturability over baselines

We perform a qualitative evaluation of Sketch2Prototype by generating 3D models with two baseline approaches. The first approach is directly passing an unprocessed sketch into Shap-E to generate the resulting mesh. The second approach is passing our sketch into ControlNet to generate 4 candidate images. We then pass each image into Shap-E and, for standardization, select the first generated mesh. For Sketch2Prototype, we also generate 4 candidate images and perform the same mesh selection process. We test our method on a phone stand design and a pen-and-coin holder. Sketch-to-text via GPT-4V, text-to-image via DALL-E 3 or ControlNet, and image-to-3D via Shap-E each take a matter of seconds, so the time difference between these three approaches is negligible.

Results are shown in Figure 5A. We can see that the meshes generated from unprocessed sketches are sparse and unmanufacturable. For ControlNet, the generated images lack diversity. In the case of the phone stand, due to the simplicity of the sketch, this process also produces unmanufacturable designs. The ControlNet generated images lack diversity due to directly matching the sketch geometry. Finally, our model generates diverse variations of each sketch. Our model also generated functional and cohesive meshes for prototyping.

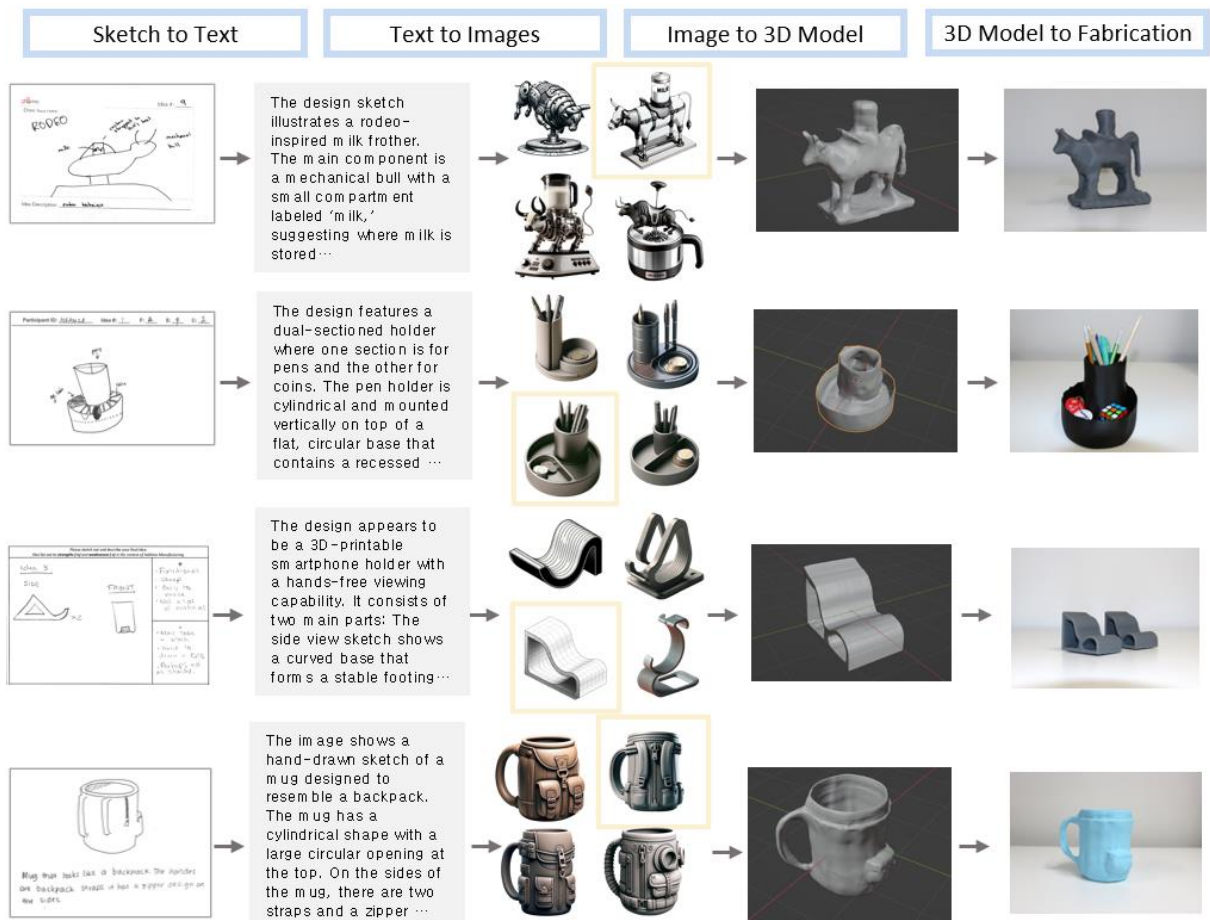


Figure 4. Examples of the full Sketch2Prototype framework for four different design types: rodeo-inspired milk frother, pen-and-coin holder, phone stand, and backpack-inspired mug

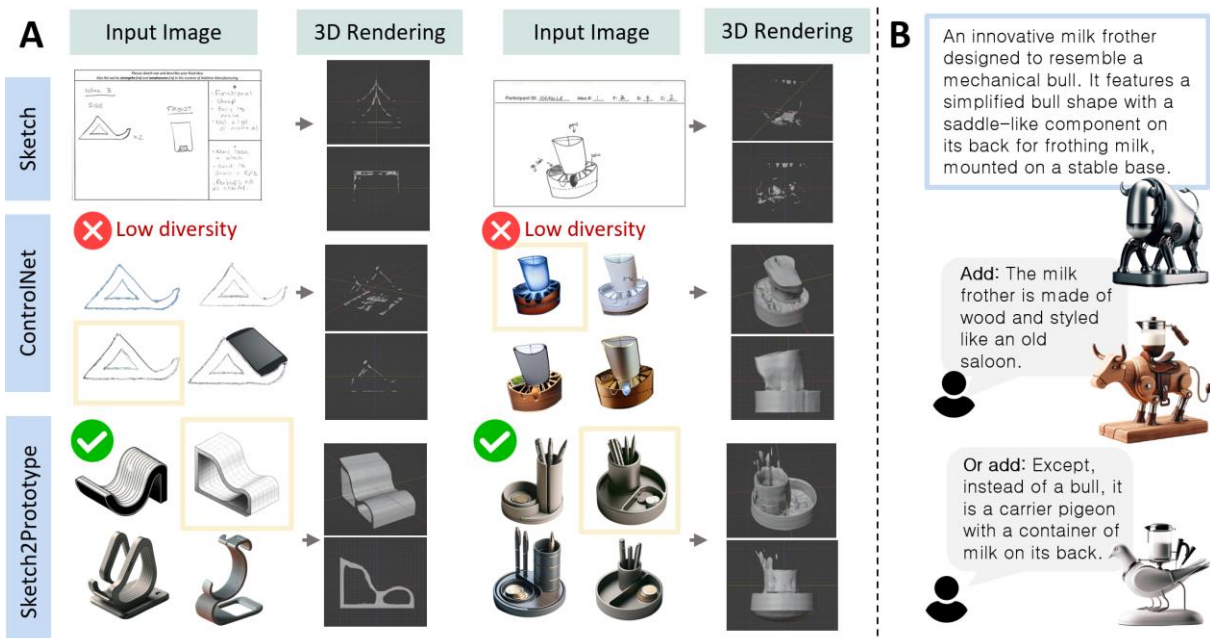


Figure 5. A: 3D models generated from varying input images: Sketch2Prototype generates more diverse and manufacturable designs. B: The text modality allows for user control. We append text to the original prompt to generate different designs

4.2. Generation and exploration via human-in-the-loop feedback

We also evaluate the controllability of Sketch2Prototype via the text modality. In Figure 5B, we add sentences to the DALL-E 3 prompt to alter the output image according to designer feedback. For example, when the original prompt is appended with "The milk frother is made of wood and styled like and old saloon," the output image changes form accordingly. The intermediate text modality thus enables users to add iterative feedback, allowing for extra user control. Hence, there is immense value in using text as an intermediary modality to help with exploration and addition of new requirements, which are difficult to update in sketch directly.

4.3. Dataset alignment and diversity

We show quantitatively that Sketch2Prototype generates images that match the original description while being diverse. We generated a synthetic dataset from a set of milk frother designs and evaluated its alignment with the original dataset and its diversity. This dataset of milk frothers contains 359 hand-drawn sketches of milk frothers drawn by unique individuals. Each milk frother design often has a brief text description of the design concept in addition to text annotations on the sketch. To generate the synthetic dataset, we provided the image to GPT-4V and incorporated the text description into prompt that asks GPT-4V to generate a DALL-E 3 prompt. For each sketch, we generate 4 images, giving us a synthetic dataset of 1,436 total images.

For our generated dataset, we show 1) that the generated images align with their sketches, and 2) that the generated images represent a diverse set, thus exploring the design space. To show our synthetic dataset aligns with the original dataset, we compute the CLIP score between each sketch in the original dataset and the provided text description. We also take the average CLIP score between the 4 images in the augmented dataset with the text description. Finally, we take the CLIP score between the sketches and the images. The results are listed in the plot on Figure 6A.

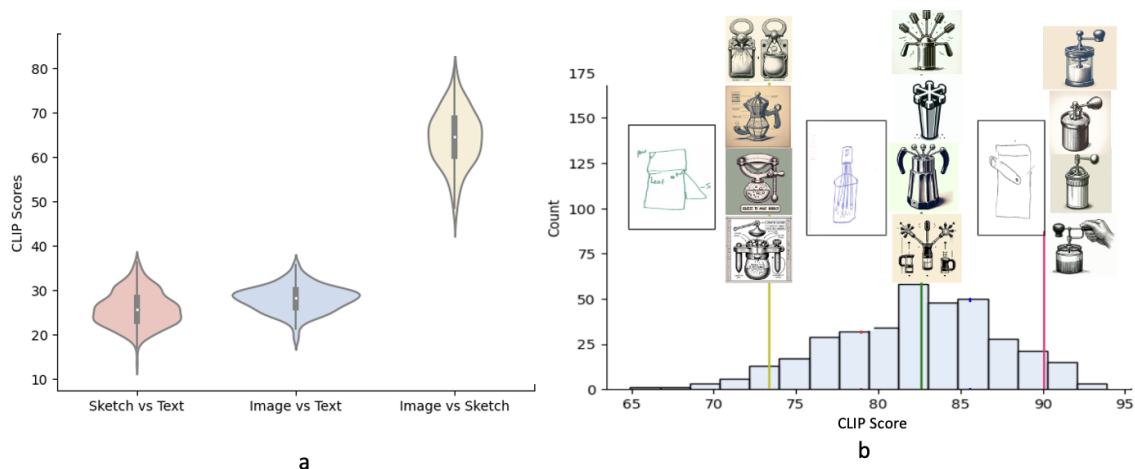


Figure 6. (a) CLIP scores between Sketch and Text; Image and Text; and Image and Sketch (b) Average pairwise CLIP score between the four generated images. A lower CLIP score indicates a more diverse set of images. Yellow, red, green, blue and pink lines correspond to 5th, 50th, and 95th percentiles respectively. The images at the very top are the original sketches

The resulting average CLIP score for the sketch vs text, image vs text and sketch vs image sets are 25.8, 28.1 and 64.4 respectively. The higher average between the image vs text and sketch vs text is expected, since GPT-4V gave a more detailed text representation for generation compared to the original sketch. Likewise, CLIP scores between two images are generally higher than CLIP scores between image and text, so it is expected that the image-sketch CLIP score is higher than the remaining two.

To measure diversity, we compute the average pairwise CLIP score for each of the 4 images in each sample. Higher CLIP scores indicate less diverse datasets. Figure 6B illustrates the distribution of these pairwise CLIP scores. At the 5th, 50th, and 95th percentiles, we show an example of a sketch and four generated images to give better context on how these scores correspond to the diversity images. As percentile increases, we observe a decrease in geometric variation and fewer components. At the 95th

percentile, the milk frother designs shown have almost identical geometries - every design are variations of a cylindrical container attached to a handle. However, designs at the 5th percentile have varied shapes and multiple components, such as handles of different curvatures, and different containers.

5. Discussion

5.1. Integrating AI into the design thinking process

In this work, we proposed a framework, Sketch2Prototype, that utilizes generative AI to rapidly generate a textual description, a diverse set of 2D images, and 3D models from a hand drawn sketch. We compare our framework to two baseline models, using sketch-alone and using ControlNet-generated images. We showed that our framework generates more diverse designs and manufacturable models than the others (Figure 5A). Sketch2Prototype also increases the breadth and depth of exploration by allowing designers to work with sketches and prototypes in parallel. We exhibit the entire framework, resulting in six fabricated prototypes from four hand-drawn sketches.

We also demonstrate how user feedback can be incorporated via the text modality. Our framework represents a design as a sketch, text, image, and 3D model. The intermediate text modality can be easily edited by a designer to add more requirements that are not present in original text. This allows iterative refinement and improvement, shown in Figure 5B. The two results shown in Figure 5 demonstrate the balance between user-control and automatic design expansion. While ControlNet allows for strict geometric adherence to an input image, this eliminates most diversity and may not be desirable when dealing with an imprecise hand-drawn sketch. On the other hand, DALL-E 3 generates very diverse designs, however iterative feedback via text may be necessary to generate a desired image.

Tools such as Sketch2Prototype can be incorporated into traditional design thinking processes, enhancing ideation, and prototype development. Sketch2Prototype enables engineers to rapidly explore the design space by expanding simple, abstract sketches into diverse images and 3D printable looks-like prototypes. Our framework allows for human-in-the-loop feedback in the text-to-image phase, and facilitates prototype development by converting these images into 3D models.

5.2. Limitations within the existing framework

We must emphasize that this is a preliminary exploration of how emerging AI tools may benefit designers via accelerated design transformation from sketch to text, images, and 3D models. The 3D models are meant to be looks-like, not functional, prototypes that can be built using additive manufacturing techniques. Further studies are necessary to determine which AI tools perform best at each step of the framework, and, in fact, define what "best" means.

The process of combining different design transformation steps means that deviance from the desired design can be introduced at every step. In the sketch-to-text phase, the user benefits from having much control via simple text editing. However, we observe a failure case in sketch-to-text then text-to-image when GPT-4V generates a text description that is deemed an "unsafe" prompt by DALL-E 3. This indicates that GPT-4V's text generation is not grounded on DALL-E 3's safety mechanisms.

Another limitation is the lack of repeatability in the text-to-image stage. Using off-the-shelf image generation tools such as DALL-E 3 means that the same prompt will not generate the same image when repeated. Furthermore, changes to the design are often global rather than local, even if the textual prompt only requests a local change.

The greatest limitation that we observe is in the image-to-3D stage. Though we utilize state of the art image-to-3D models, design details are often lost at this stage, and 3D models generated using image-to-3D are often non-smooth, fragmented, and sometimes non-manufacturable. As a result, postprocessing is required either to smooth surfaces, fill holes, or remove unmanufacturable parts. Even with postprocessing, image-to-3D models tend to create uneven surfaces or holes, which indicate that NeRFs, while strong for 3D visualization purposes, may not be ideal for 3D printing. Due to the postprocessing step, the time taken for Sketch2Prototype increases.

Finally, aside from manufacturability problems, there may be a lack of control over the final mesh. For earlier steps in the framework, users can edit intermediate representations to better control the product,

such as editing descriptions during sketch-to-text or inpainting during text-to-image. The only control users have over a mesh is the postprocessing stage, which is limited to removal or minor edits to surfaces. Enabling users to "inpaint" 3D meshes would give significantly more flexibility over their final product.

6. Future work

Future research could explore the application of this framework in more complex design scenarios, such as multi-component systems or intricate structures. These tasks are challenging since models would need to identify distinct parts, understand part-interfaces, and ensure compatibility. This challenge becomes even harder when dealing with sketches that have internal components, as internal components need to have correct proportions with respect to their container in order to fit, and we note that the image-to-3D models we use did not capture internal components well.

Current research on image-to-3D models is often concerned with synthesizing 3D images from objects for visualization purposes, but making these meshes functional is much harder. This work reveals that NeRFs may not be an ideal candidate for representing manufacturable prototypes. An area for exploration may be to create a new representation of meshes specifically for 3D printing purposes.

This work aims to demonstrate how existing AI tools enable transformation from sketch to text, image, and 3D modalities, which can enhance design space exploration. Sketch2Prototype should assist designers in efficiently ideating with different modalities. To this end, future work may include incorporating user-centered design principles in the Sketch2Prototype framework, focusing on intuitive interfaces and user feedback.

7. Conclusion

We demonstrate a framework to convert sketches into fabricated prototypes via intermediate steps: sketch-to-text, text-to-image, and image-to-3D. We show that our framework enhances design space exploration by generating a set of diverse 2D images and 3D models from a single sketch, and by enabling designers to work with sketches and prototypes in parallel. We find that using text as an intermediate modality allows for iterative user feedback and enhanced user control. Furthermore, text-to-image-to-3D generates more diverse and manufacturable 3D models than sketch-to-3D baselines. However, manufacturability is still a limitation of current image-to-3D models. The Sketch2Prototype framework gains potential as each step is actively worked on by the machine learning community.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 2231254 and the NSF Graduate Research Fellowship.

References

- Bao, Q., Faas, D., & Yang, M. (2018). Interplay of sketching & prototyping in early stage product design. *Int. J. of Design Creativity and Innovation*, 6(3-4), 146-168. <https://dx.doi.org/10.1080/21650349.2018.1429318>.
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., & Hedman, P. (2021). Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. *CoRR*, abs/2111.12077.
- Camburn, B., Viswanathan, V., Linsey, J., Anderson, D., Jensen, D., Crawford, R., Otto, K., & Wood, K. (2017). Design prototyping methods: State of the art in strategies, techniques, and guidelines. *Design Sci.*, 3, E13. <https://dx.doi.org/10.1017/dsj.2017.10>.
- Corbett, J., & Crookall, J.R. (1986). Design for Economic Manufacture. *CIRP Ann.*, 35(1), 93-97. [https://doi.org/10.1016/S0007-8506\(07\)61846-0](https://doi.org/10.1016/S0007-8506(07)61846-0).
- Das, M., & Yang, M.C. (2022). Assessing Early Stage Design Sketches and Reflections on Prototyping. *ASME J. Mech. Des.*, 144(4): 041403. <https://dx.doi.org/10.1115/1.4053463>.
- Edwards, K.M., Peng, A., Miller, S.R., & Ahmed, F. (2022). If a Picture is Worth 1000 Words, Is a Word Worth 1000 Features for Design Metric Estimation?. *ASME J. Mech. Des.*, 144(4): 041402. <https://doi.org/10.1115/1.4053130>.
- Edwards, K.M., Addala, V.L., & Ahmed, F. (2021). Design Form and Function Prediction From a Single Image. In *ASME IDETC-CIE. Virtual*, Online. <https://dx.doi.org/10.1115/DETC2021-71853>.
- Elverum, C.W., Welo, T., & Tronvoll, S. (2016). Prototyping in New Product Development: Strategy Considerations. *Procedia CIRP*, 50, 117-122. <https://dx.doi.org/10.1016/j.procir.2016.05.010>

- Goldschmidt, G. (2014). Modeling the Role of Sketching in Design Idea Generation. In: Chakrabarti, A., Blessing, L. (eds) *Anth. of Theory and Models of Des.* Springer, London.
- Han, J., Forbes, H., & Schaefer, D. (2021). An exploration of how creativity, functionality, and aesthetics are related in design. *Res Eng Design*, 32, 289–307. <https://dx.doi.org/10.1007/s00163-021-00366-9>.
- Hansen, C.A., & Özkil, A.G. (2020). From Idea to Production: A Retrospective and Longitudinal Case Study of Prototypes and Prototyping Strategies. *ASME J. Mech. Des.*, 142(3): 031115. <https://doi.org/10.1115/1.4045385>.
- Hessel, J., et al. (2022). CLIPScore: A Reference-Free Evaluation Metric for Image Captioning. arXiv:2104.08718, arXiv, 23 Mar. 2022. arXiv.org, <http://arxiv.org/abs/2104.08718>.
- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising Diffusion Probabilistic Models. CoRR, abs/2006.11239.
- Jun, H., & Nichol, A. (2023). Shap-E: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463.
- Kerbl, B., Kopanas, G., Leimkuehler, T., & Drettakis, G. (2023). 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4), 139:1-139:14. <https://dx.doi.org/10.1145/3592433>.
- Lauff, C., Menold, J., & Wood, K. (2019). Prototyping Canvas: Design Tool for Planning Purposeful Prototypes. *Proc. of the Design Soc.: Int. Conf. on Eng. Design*, 1(1), 1563-1572. <https://dx.doi.org/10.1017/dsi.2019.162>.
- Melas-Kyriazi, L., Rupprecht, C., Laina, I., & Vedaldi, A. (2023). RealFusion: 360{deg} Reconstruction of Any Object from a Single Image. CVPR, 23 February.
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., & Ng, R. (2021). NeRF: representing scenes as neural radiance fields for view synthesis. *Commun. ACM*, 65(1), 99–106. <https://dx.doi.org/10.1145/3503250>.
- Murugappan, S., Piya, C., Yang, M.C., & Ramani, K. (2017). FEAsy: A Sketch-Based Tool for Finite Element Analysis. *ASME J. Comput. Inf. Sci. Eng.*, 17(3): 031009. <https://doi.org/10.1115/1.4034387>.
- Neeley, L., Lim, K., Zhu, A., & Yang, M.C. (2013). Building Fast to Think Faster: Exploiting Rapid Prototyping to Accelerate Ideation During Early Stage Design. *ASME Int. Design Eng. Tech. Conf.*, Portland, OR, 2013. <https://dx.doi.org/10.1115/DETC2013-12635>.
- OpenAI. (2023). GPT-4 Technical Report. arXiv:2303.08774, arXiv, 27 Mar. 2023. arXiv.org, <http://arxiv.org/abs/2303.08774>.
- Pahl, G., Beitz, W., Feldhusen, J., & Grote, K.-H. (2007). *Engineering Design: A Systematic Approach*. Springer London. <https://dx.doi.org/10.1007/978-1-84628-319-2>.
- Picard, C., Edwards, K.M., Doris, A.C., Man, B., Giannone, G., Alam, M.F., & Ahmed, F. (2023). From Concept to Manufacturing: Evaluating Vision-Language Models for Engineering Design. arXiv preprint arXiv:2311.12668.
- Poole, B., Jain, A., Barron, J.T., & Mildenhall, B. (2022). DreamFusion: Text-to-3D using 2D Diffusion. arXiv, 29 September.
- Radford, A., et al. (2021). Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020, arXiv, 26 Feb. 2021. arXiv.org, <http://arxiv.org/abs/2103.00020>.
- Saharia, C., et al. (2022). Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. arXiv:2205.11487, arXiv, 23 May 2022. arXiv.org.
- Schmidt, L., Hernandez, N., & Ruocco, A. (2012). Research on encouraging sketching in engineering design. *AI EDAM*, 26(3), 303-315. <https://dx.doi.org/10.1017/S0890060412000169>.
- Singh, A., Hu, R., Goswami, V., Couairon, G., Galuba, W., Rohrbach, M., & Kiela, D. (2022). FLAVA: A Foundational Language And Vision Alignment Model. CVPR, 29 March.
- Song, B., Miller, S., & Ahmed, F. (2023). Attention-Enhanced Multimodal Learning for Conceptual Design Evaluations. *ASME J. Mech. Des.*, 145(4): 041410. <https://dx.doi.org/10.1115/1.4056669>.
- Tang, J., Ren, J., Zhou, H., Liu, Z., & Zeng, G. (2023). DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. arXiv, 28 September.
- Toh, C.A., & Miller, S.R. (2016). Choosing creativity: the role of individual risk and ambiguity aversion on creative concept selection in engineering design. *Res. Eng. Design*, 27, 195–219. <https://doi.org/10.1007/s00163-015-0212-1>.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., et al. (2023). LLaMA: Open and Efficient Foundation Language Models. arXiv, 27 February.
- Ulrich, K.T., Eppinger, S.D., & Yang, M.C. (2020). *Product Design and Development*. New York, NY: McGraw-Hill Education.
- Yu, A., Ye, V., Tancik, M., & Kanazawa, A. (2021). Pixelnerf: Neural radiance fields from one or few Images. CVPR 2021