

Introducing 3D Sketching to Overcome Challenges of View-Consistency and Progressive Development in 2D Generative AI-Based Car Exterior Design

Seung-Jun Lee
Department of Industrial Design
KAIST
Daejeon, Republic of Korea
seung-jun.lee@kaist.ac.kr

Joon Hyub Lee
DRB-KAIST SketchTheFuture
Research Center
KAIST
Daejeon, Republic of Korea
joonhyub.lee@kaist.ac.kr

Seok-Hyung Bae
Department of Industrial Design
KAIST
Daejeon, Republic of Korea
seokhyung.bae@kaist.ac.kr

Abstract

Recent advances in 2D generative AI are beginning to find applications in highly specialized fields, such as car exterior design. However, the current 2D-centric approach has several limitations: each viewpoint requires a new sketch; maintaining consistency across different viewpoints is challenging; steering design development in the desired direction can be difficult. To address these limitations, we propose a novel design workflow that integrates 3D sketching with 2D generative AI for car exterior design. This workflow enables car designers to seamlessly transition between expressive 3D sketching, detailed 2D drawing, and realistic 2D generation, facilitating view-consistent and progressive design development. We conducted an in-depth user test with a professional car designer, who used our system to produce car exterior concepts for all major body types, demonstrating its potential usefulness during the early stages of car design.

CCS Concepts

• Human-centered computing → Interaction techniques.

Keywords

Car design; 3D sketching; 2D generative AI

ACM Reference Format:

Seung-Jun Lee, Joon Hyub Lee, and Seok-Hyung Bae. 2025. Introducing 3D Sketching to Overcome Challenges of View-Consistency and Progressive Development in 2D Generative AI-Based Car Exterior Design. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3706599.3719731>

1 INTRODUCTION

Thanks to rapid advances in 2D generative AI based on diffusion models, car designers can now easily obtain realistic renderings from minimal sketches that convey their design intent. This capability allows designers to quickly explore a wide range of possibilities

Permission to make digital or hard copies of all or part 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).

CHI EA '25, Yokohama, Japan

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1395-8/25/04

<https://doi.org/10.1145/3706599.3719731>




Figure 1: We propose a novel car exterior design workflow that integrates 3D sketching and generative AI in a complementary manner. Within this workflow, car designers can seamlessly transition between expressive 3D sketching, detailed 2D drawing, and realistic 2D generation, enabling view-consistent and progressive design development.

and develop them into concrete concepts during the early stages of the design process.

However, there are several limitations to applying 2D generative AI in professional car design practice:

- Once a satisfactory rendering is achieved from one viewpoint, designers must create entirely new sketches to visualize the design from other viewpoints, requiring significant time and effort.
- Images generated from sketches of different viewpoints often lack view-consistency, and as a result, the set of images may not describe the same 3D shape.
- As designers repeat the cycle of sketching and generation, design intent may be lost in translation, making it difficult to accumulate intermediate results and steer design development in the desired direction.

We believe that leveraging 3D sketching can help overcome these limitations. It allows designers to intuitively express their ideas as 3D curves through minimal gestures, similar to those used in 2D sketching. Once created, the 3D sketch can be viewed from any angle and serve as an enduring structural backbone onto which new ideas can be progressively added.

In this paper, we propose a novel car design workflow that integrates the strengths of 3D sketching and 2D generative AI in a complementary manner. This workflow enables designers to seamlessly transition between 3D sketching, 2D drawing, and 2D generation, ensuring view-consistent and progressive development

of car exterior designs across multiple viewpoints (Figure 1). We conducted an in-depth user test with an experienced car designer for approximately 14 hours to investigate the potential benefits this workflow could bring to early-stage design practices.

2 RELATED WORK

In this section, we highlight the importance of sketching in the early stages of car design and introduce 3D sketching techniques suitable for this context. Furthermore, we outline advances in diffusion-based 2D and 3D generative models and discuss efforts to apply these models to car design, as well as their limitations, emphasizing the need for our approach.

2.1 Car Design Process

In the early concept design stage, designers define the car’s target market segment, functional goals, and packaging, including its size, proportions, and layout [13]. Based on these specifications, they engage in intensive and competitive 2D sketching [11]. Designers rapidly generate numerous ideas through concise line drawings, some of which are selected and realistically rendered for communication with colleagues [4]. Through iterative design reviews, refined concepts are developed into scale or full-size clay models, and subsequently into sophisticated engineering CAD models for production [13]. This study aims to support the concept design stage, where designers must rapidly explore, communicate, and refine diverse possibilities within a short timeframe.

2.2 3D Sketching

3D sketching is an interactive technique that allows designers to express form ideas as 3D curves using pen input, much like drawing on a digital tablet, facilitating rapid exploration, communication, and refinement of many ideas at a level of fidelity appropriate for the early design stage [1]. The resulting 3D sketch can not only be viewed from any angle, but also reviewed at real scale in VR [10], or used as reference data to accelerate subsequent modeling and engineering stages [1, 8, 9].

One commonly used 3D sketching method involves projecting 2D curves onto predefined 3D planes, resulting in 3D planar curves [6–9, 15, 25]. While simple and powerful, this method may be inadequate for handling the more sophisticated and nuanced forms of cars [4]. Therefore, we also incorporate additional techniques for producing 3D spatial curves [1], such as the orthographically extruded surface method, the single-view symmetric epipolar method, and the two-view epipolar method, as a comprehensive set of tools for expressing car exterior designs.

2.3 Diffusion-Based 2D and 3D Generation

Diffusion models have achieved groundbreaking advances in high-quality 2D image generation through iterative noise addition and removal [21]. Researchers have improved and optimized this approach [5, 22], enabling realistic image generation at high speed. This progress has given rise to models that generate 2D images from text prompts [17, 19, 20], as well as models that generate 2D images from 2D sketches [14, 24, 26].

While 3D mesh can also be generated from text prompts [12, 16] and 2D images [3], their results are typically of lower quality compared to 2D image generation. Since most early-stage design communication and decision-making in practice rely on high-quality 2D images [4], we focus on a workflow centered around 2D image generation rather than 3D mesh generation.

2.4 Generative AI-Based Car Design

Recently, commercial services, such as Vizcom [23] and Optic [2], have made high-quality 2D car rendering generation from 2D sketches and text prompts readily available for professional car designers. However, designing from multiple viewpoints still requires creating new 2D sketches from scratch, and the resulting images may not represent a consistent 3D form.

Specifically, while Vizcom offers 3D mesh generation from 2D images, it produces incompatible meshes from images of the same car viewed from different viewpoints, with each mesh failing to incorporate visual features not visible from that particular viewpoint. Optic enables the use of 3D meshes as references to help enforce consistency across generated 2D images but requires additional 3D CAD modeling alongside 2D sketching. Instead, this study leverages 3D sketches, which can be quickly created and viewed from any angle, serving as a structural backbone to ensure consistency across 2D images.

3 SYSTEM

We propose a novel interactive system for a car exterior design workflow that seamlessly integrates 3D sketching, 2D drawing, and 2D generation in the early design stage (Figure 1). In this workflow, designers can progressively develop their concepts while maintaining a high level of view consistency through the following steps:

- Quickly sketch essential 3D curves to express initial ideas (Figures 2a, 3, 4).
- Add 2D details on top of the 3D sketch rendered as a 2D line art underlay from a desired viewpoint to flesh out design intent (Figures 2b, 5a).
- Generate realistic 2D images with AI to explore design options and gain inspiration (Figure 5b, c).
- Update the 3D sketch backbone with new 3D curves based on the generated outcomes to accumulate promising ideas, and repeat these steps as many times as needed.

3.1 2D and 3D Workspace

We provide designers with a 3D workspace for 3D sketching and multiple 2D workspaces for 2D drawing and 2D generation, which they can create from any viewpoint. To facilitate seamless transitions between these workspaces, we provide three widgets: thumbnails, reference images, and flagsticks (Figure 2).

In addition, since car designers often begin ideation by referencing existing designs [18], our system allows importing 3D car models as translucent, monochrome templates.

3.2 3D Sketching

Designers can utilize one 3D planar curve sketching method and three 3D spatial curve sketching methods to create essential 3D curves that express car exterior designs. Designers can erase these curves partially or entirely, and the system automatically mirrors all curves.

- 3D planar curve sketching method: the orthographic plane method (Figure 3a).
- 3D spatial curve sketching methods: the orthographically extruded surface method (Figure 3b, c), the single-view symmetric epipolar method (Figure 4a, b), and the two-view epipolar method (Figure 4a, c).

3.3 2D Drawing

Designers can capture 3D sketches from desired viewpoints and use them as underlays for 2D drawings, where certain design intent, such as defining silhouettes, adding details, and shading surfaces, can be expressed more efficiently (Figure 5a).

Professional car designers develop deep familiarity and expertise with specialized 2D image editing tools through years of practice. Therefore, instead of duplicating these tools, our system provides seamless integration, allowing designers to export images, modify them in their preferred software, and easily re-import them automatically.

3.4 2D Generation

Designers can generate multiple 2D renderings using a 2D drawing as input, effortlessly previewing how their early ideas could evolve in many different directions (Figure 5b, c). These high-quality 2D renderings can then serve as a foundation for further 2D drawing and 3D sketching, facilitating iterative design development.

Since many 2D generative AI models with various features and tradeoffs are being actively developed and released, our system allows designers to choose and replace models according to their specific needs.

4 IMPLEMENTATION

We implemented our system using the Unity 3D game engine. For 2D drawing, we selected Photoshop, the industry-standard application, and used Unity’s PythonRunner and the comtypes Python package to automate file exchanges. For 2D generation, we chose Vizcom, a state-of-the-art commercial web service specializing in industrial and car design, and used TCP for data transfers.

The system was executed on a Lenovo Legion 5i gaming laptop with Windows 11 OS, Intel Core i9-14900HX CPU, 32 GB of RAM, and Nvidia GeForce RX 4070 GPU, paired with a Wacom Cintiq Pro 24 Touch digital tablet supporting both multi-touch and pen input capabilities.




Figure 2: 2D and 3D workspace. The designer can (a) perform 3D sketching on top of a car body template within the 3D workspace, and (b) create 2D workspaces from any desired viewpoint. The system organizes all 2D outputs as thumbnails (outlined in red), which designers can tap to enlarge or drag into the overlay to create reference images (outlined in green). When the designer rotates the view and exits to the 3D workspace, (c) flagsticks (outlined in orange) appear, marking the positions of the 2D workspaces. Tapping on a reference image or a flagstick reopens the corresponding 2D workspace.




Figure 3: 3D sketching with orthographic plane method and orthographically extruded surface method. The designer can double-tap a grid plane to access an orthographic view and (a) draw on the grid plane (in this case, the center plane) to create a 3D planar curve (in this case, from the side view). Upon rotating away from the view, (b) the designer can tap a planar curve to create a sketch surface extruded along the normal direction of the grid plane, and (c) draw a spatial curve on it.




Figure 4: 3D sketching with single-view symmetric epipolar method and two-view epipolar method. The designer can (a) draw a curve to create an epipolar sketch surface, then either (b) draw a symmetric pair from the same view, or (c) draw the initial curve again from another view to create 3D spatial curves.




Figure 5: 2D drawing and 2D generation. After creating a 2D workspace, (a) the system performs hidden-line removal, renders the 3D sketch as 2D line art, and automatically launches an external 2D editing tool (in this case, Photoshop) for 2D drawing. After the 2D drawing is completed, (b) the designer can configure parameters including text prompts, drawing influence levels, and the number of 2D images to generate. After the 2D generation is completed, resulting in a row of thumbnails, (c) the designer can drag-and-drop a thumbnail to create a reference image for further use.

5 IN-DEPTH USER TEST

We conducted an in-depth user test to investigate how our proposed workflow could benefit car exterior design practice. We invited a veteran car designer with extensive industry experience to create new design concepts across all major car body types using our implemented system. We recorded the designer's usage of the system and held a follow-up interview for detailed insights.

5.1 Participant

The participant (male, 41 years old) was a professional car designer with 15 years of work experience. He directed over 20 concept and production projects at major international car manufacturers and design studios and has received prestigious international design awards for his work, including the Red Dot Design Award, the iF Design Award, the Good Design Award, and the Spark Award.

5.2 Task

The designer was asked to create new design concepts across all 13 major car body types [13] and produce a set of rendered images for each car from three key viewpoints (front three-quarter, side, and rear three-quarter), which are most commonly used in car design practice for effective communication and decision-making. We provided the existing car models the designer had requested as templates for 3D sketching (Table 1).

5.3 Procedure

The in-depth user test comprised three sequential sessions:

1. Tutorial session: The designer was instructed on how to use the system and allowed to practice freely until he was fully familiar with it.
2. Task session: The designer performed tasks for 13 car body types in randomized order.
3. Debriefing session: The designer and the experimenter revisited the design process by reviewing time-lapses of the recorded videos together and discussed the designer's intent and results at each step of the process.

5.4 Measurement

We measured the frequency, duration, and order of each activity, including 3D sketching, 2D drawing, 2D generation, and idling. We also collected the number of 3D sketches, 2D drawings, and 2D generations, as well as the number of 3D curves created and the sketching methods used. Throughout the task session, we positioned a camera behind the designer's shoulder and recorded the screen and hand interactions.

5.5 Result

The designer completed all tasks in 13 hours and 55 minutes, creating 39 renderings of 13 car exterior concepts, a selection of which are presented in this paper (Figures 6-10). We processed the system usage logs and calculated aggregate numbers (Table 1). In addition, we reconstructed the process by which designs were progressively developed and visualized it in a timeline format (Figure 11).




Figure 6: The 3D sketch and rendered images produced by the designer from the front three-quarter, side, and rear three-quarter views of the one-box body type (total time spent (h:mm): 1:18; number of 3D sketches (3DS): 3; number of 2D drawings (2DD): 9; number of 2D generations (2DG): 48).




Figure 7: Two-box type (0:57; 3DS: 3; 2DD: 7; 2DG: 29).




Figure 8: Coupe type (1:09; 3DS: 3; 2DD: 8; 2DG: 31).




Figure 9: SUV type (0:45; 3DS: 4; 2DD: 9; 2DG: 28).




Figure 10: Pickup truck type (1:27; 3DS: 3; 2DD: 7; 2DG: 37).

6 DISCUSSION

Based on the system usage log collected during the task session and the feedback from the debriefing session, we discuss how the designer achieved satisfactory results by utilizing 3D sketching, 2D drawing, and 2D generation synergistically, taking advantage of our integrated and flexible workflow.

The designer rapidly produced view-consistent results that met the standards required in professional practice. He quickly learned our system and successfully created sets of 2D renderings with high view-consistency and completeness within an average of 1 hour and 4 minutes per car (min: 0:37, max: 1:27). He assessed his work as being of “quality suitable for idea communication in professional car design practice” and expressed satisfaction with the results. He estimated that achieving similar results using traditional workflows involving only 2D drawing and 2D generation would take 4 to 5 times longer.

The designer fluently expressed 3D ideas using the provided 3D sketching methods. He spent an average of 29.1% of his time on 3D sketching, frequently utilizing the orthographically extruded surface method (63.7%), the two-view epipolar method (17.7%), the orthographic plane method (16.1%), and the single-view symmetric epipolar method (2.5%), in that order.

- **Early stage.** He used the orthographic plane method to establish the overall proportions of the cars. In particular, he began by sketching the side views, noting that “designers usually start with side views because they clearly show cars without any perspective distortion.”
- **Middle stage.** Once the proportions of the cars were set, he applied the orthographically extruded surface method to develop 3D planar curves created with the orthographic plane method into 3D spatial curves. At the same time, obsolete 3D planar curves were erased.
- **Late Stage.** From the point when rough 3D shapes of the cars were visualized, he primarily utilized the epipolar methods, noting that “once the cars’ overall proportions were set, the epipolar

methods allowed me to sketch freely and capture the intended impression, much like drawing in a sketchbook.”

The designer complemented 3D sketching and 2D generation with 2D drawing. He spent an average of 49.0% of his time on 2D drawing. Specifically, about a quarter of this time (11.2% of the total time) was spent creating input for 2D generation using 3D sketches as underlays, and three quarters (37.8% of the total time) were spent fine-tuning the outputs of 2D generation.

- **Creating input for 2D generation with minimal 2D drawings on top of 3D sketches.** Instead of directly using barebone 3D sketches as input for 2D generation, he found that enhancing the underlays with minimal 2D drawings resulted in higher-quality outputs that better aligned with his intentions. For instance, on top of the 3D sketch underlays, he added missing silhouette lines, details such as headlights, grilles, and door lines, and simple shading. Explaining the benefits of this complementary approach, he noted that “since most of the shape was already there on the canvas when I started 2D drawing, I could quickly try many ideas.”
- **Fine-tuning 2D generation outputs for the final rendering sets.** After generating and selecting desired images for the three viewpoints of each car, he refined them in Photoshop, comparing them against each other to create a cohesive final rendering set. He noted that “the AI-generated images were generally consistent and high-quality, but they often deviated from my intentions in the details. It still took considerable work to stamp out these differences.” He hoped that as generative AIs for 2D images improve in accuracy as well as speed, the time required for such adjustments would be significantly reduced in the future.

The designer progressively developed his designs by updating 3D sketches based on 2D generation. He spent an average of 16.8% of his time on 2D generation, producing 32.2 images per car. Rather than a linear progression, his system usage pattern revealed multiple iterative cycles, often returning to 3D sketching after 2D generation (Figure 11).

Body type [13]	Template	Time spent (h:mm (%))					Output (#)				3D curve (# (%))			
		3DS	2DD	2DG	Idle	Total	3DS	2DD	2DG	OP	OES	SVSE	TVE	Total
One-box	Citroen Ami '21	0:19 (24.3)	0:45 (57.7)	0:12 (15.4)	0:02 (2.6)	1:18 (100)	3	9	48	79 (16.4)	332 (69.0)	0 (0)	70 (14.6)	481 (100)
Two-box	Volkswagen Golf '20	0:16 (28.1)	0:30 (52.6)	0:09 (15.8)	0:02 (3.5)	0:57 (100)	3	7	29	77 (21.2)	228 (62.8)	0 (0)	58 (16.0)	363 (100)
Three-box	BMW 5 Series '21	0:11 (19.3)	0:34 (59.7)	0:10 (17.5)	0:02 (3.5)	0:57 (100)	3	6	40	68 (18.4)	182 (49.2)	12 (3.2)	108 (29.2)	370 (100)
Limousine	Rolls Royce Phantom '23	0:09 (24.3)	0:22 (59.5)	0:04 (10.8)	0:02 (5.4)	0:37 (100)	1	6	24	52 (30.2)	108 (62.8)	2 (1.2)	10 (5.8)	172 (100)
Wagon	Volvo V60 '19	0:14 (19.7)	0:48 (67.6)	0:07 (9.9)	0:02 (2.8)	1:11 (100)	3	6	24	55 (16.9)	210 (64.6)	28 (8.6)	32 (9.9)	325 (100)
Coupe	Mercedes-Benz AMG GT '20	0:18 (26.1)	0:34 (49.3)	0:12 (17.4)	0:05 (7.2)	1:09 (100)	3	8	31	59 (12.6)	330 (70.3)	28 (6.0)	52 (11.1)	469 (100)
Fastback	Audi A7 Sportback '18	0:15 (22.1)	0:41 (60.3)	0:10 (14.7)	0:02 (2.9)	1:08 (100)	3	6	30	49 (20.0)	144 (58.8)	18 (7.3)	34 (13.9)	245 (100)
Mid-engine supercar	McLaren GT '20	0:22 (29.7)	0:14 (18.9)	0:35 (47.3)	0:03 (4.1)	1:14 (100)	3	6	40	75 (10.3)	460 (63.1)	2 (0.3)	192 (26.3)	729 (100)
Roadster	Alfa Romeo 4C Spider '15	0:28 (41.2)	0:31 (45.6)	0:06 (8.8)	0:03 (4.4)	1:08 (100)	3	6	25	54 (9.8)	372 (67.6)	22 (4.0)	102 (18.6)	550 (100)
SUV	Kia Sportage '23	0:18 (40.0)	0:16 (35.5)	0:08 (17.8)	0:03 (6.7)	0:45 (100)	4	9	28	41 (9.5)	154 (35.7)	0 (0)	236 (54.8)	431 (100)
Pickup truck	Ford F-150 '24	0:21 (24.1)	0:38 (43.7)	0:18 (20.7)	0:10 (11.5)	1:27 (100)	3	7	37	114 (15.8)	586 (81.2)	0 (0)	22 (3.0)	722 (100)
Minivan	Volkswagen ID. Buzz '23	0:22 (33.3)	0:36 (54.6)	0:06 (9.1)	0:02 (3.0)	1:06 (100)	3	6	29	112 (15.6)	542 (75.7)	2 (0.3)	60 (8.4)	716 (100)
Commercial van	Mercedes-Benz Sprinter '18	0:28 (46.7)	0:19 (31.7)	0:08 (13.3)	0:05 (8.3)	1:00 (100)	4	8	34	67 (13.1)	342 (66.9)	6 (1.2)	96 (18.8)	511 (100)
Average		0:18 (29.1)	0:31 (49.0)	0:11 (16.8)	0:03 (5.1)	1:04 (100)	3.0	6.9	32.2	69.4 (16.1)	306.9 (63.7)	9.2 (2.5)	82.5 (17.7)	468.0 (100)

Table 1: Time spent (3DS: 3D sketching, 2DD: 2D drawing, 2DG: 2D generation), number of outputs, and number of 3D curves created by method (OP: orthographic plane method, OES: orthographically extruded surface method, SVSE: single-view symmetric epipolar method, TVE: two-view epipolar method), organized by car body type and corresponding template.

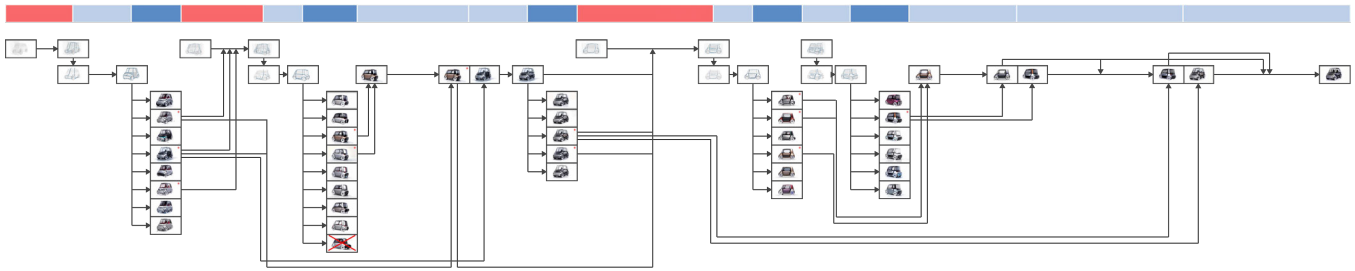


Figure 11: Reconstructed timeline of the design process for the one-box body type exterior concept. The top bar indicates activities (red: 3D sketching, light blue: 2D drawing, dark blue: 2D generation), with lengths proportional to the time spent. The first row of thumbnails shows 3D sketches, the second row shows 2D drawings, and the third and subsequent rows show 2D generations. Arrows indicate how each work influenced subsequent works.

- **Drawing inspiration from multiple viewpoints.** He explained that he often simultaneously referred to images generated from different viewpoints when updating his 3D sketch. He noted that “by glancing at these images as a whole, I could get a strong sense of the overall mood and volume the car should achieve,” adding that repeatedly revising his 3D curves to appear aesthetically pleasing from these different viewpoints helped him create “a well-rounded and balanced design.”
- **Incorporating feasibility into the ideation process.** Explaining the benefits of AI-assisted previsualization, he noted that “the realistic images generated by AI showed me what my ideas would look like if they were to become an actual car.” Moreover, regarding some pitfalls that occur when 2D and 3D processes are disconnected, he noted that “car designers often use exaggerated perspectives in 2D sketches. But if the distortion is too much, a lot of the original intentions will be gone when translated into 3D models.” Having experienced organically alternating between 2D and 3D, he claimed that “this workflow has a feasibility check built into it.”

The designer’s productivity was boosted by the tight system integration. He switched between the activities of 3D sketching, 2D drawing, and 2D generation an average of 11.0 times per car, amounting to one transition every 5.8 minutes.

- **Seamlessly transitioning between representations.** He noted that “switching modes was easy, so I could decide whether to work in 2D or 3D and make the jump without a second thought.” Additionally, with 3D sketching, Photoshop, and Vizcom automatically exchanging data in the background, he “didn’t have to manage files or windows, and could stay in the flow without losing context.”
- **Leveraging reference images for continuity and connectivity.** He collected an average of 7.2 reference images per car and noted that “having all the results gathered on the same screen assured me that the design wasn’t getting lost and that I was gradually building toward something.” He added that, over time, the spatially arranged reference images evolved from being “simple visual reminders” into “a hub connecting many ideas and their different representations.”

7 CONCLUSION & FUTURE WORK

This study proposed a novel car exterior design workflow that introduces 3D sketching to 2D generative AI-based design process to overcome limitations regarding view-consistency and progressive development. The workflow enables designers to sketch 3D curves using four different methods and organically integrates 2D and 3D workspaces through thumbnails, reference images, and flagsticks, facilitating seamless use of commercial 2D editing tools such as Photoshop and 2D image generation models such as Vizcom.

In an in-depth user test, a veteran car designer with 15 years of industry experience used our system for approximately 14 hours to create new exterior design concepts for all 13 major car body types. The outcome included 13 original 3D sketches and 39 high-quality 2D renderings. By repeatedly performing 3D sketching, 2D drawing, and 2D generation, the designer progressively developed view-consistent designs, accumulating design ideas into an evolving set of essential 3D sketch curves.

We believe that in the era of generative AI, designers’ sketching will emerge as a key tool for driving both creativity and productivity. In particular, the unique capability of 3D sketching to produce limitless 2D images from any viewpoint may become even more prominent. In addition to conducting a user study involving more professional car designers, future work could explore generating high-quality 3D meshes from 3D sketches in real time, and reviewing and refining them at real scale in a collaborative VR environment.

Acknowledgments

This research was supported by the DRB-KAIST SketchTheFuture Research Center and the KAIST Convergence Research Institute Operation Program. We thank Jeongche Yoon for the in-depth user test, and Siripon Sutthiwanna and Sang-Hyun Lee for their technical support.

References

- [1] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. I Love Sketch: as-natural-as-possible sketching system for creating 3D curve models. In *Proc. UIST '08*. 151–160.
- [2] Udin BV. 2024. *Optic*.
- [3] Minglin Chen, Weihao Yuan, Yukun Wang, Zhe Sheng, Yisheng He, Zilong Dong, Liefeng Bo, and Yulan Guo. 2024. Sketch2NeRF: multi-view sketch-guided text-to-3D generation. arXiv:2401.14257 [cs.CV] <https://arxiv.org/abs/2401.14257>
- [4] Fabio Filippini and Gabriele Ferraresi. 2021. *Curve: 15 Lezioni sul Car Design*. Rizzoli Lizard.

- [5] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Proc. NuerIPS '20*, Vol. 33. 6840–6851.
- [6] Kiia Kallio. 2005. 3D6B editor: projective 3D sketching with line-based rendering. In *Proc. SBIM '05*. 73–79.
- [7] Yongkwan Kim, Sang-Gyun An, Joon Hyub Lee, and Seok-Hyung Bae. 2018. Agile 3D sketching with air scaffolding. In *Proc. CHI '18*. Article 238, 12 pages.
- [8] Yongkwan Kim and Seok-Hyung Bae. 2016. SketchingWithHands: 3D sketching handheld products with first-person hand posture. In *Proc. UIST '16*. 797–808.
- [9] Joon Hyub Lee, Hanbit Kim, and Seok-Hyung Bae. 2022. Rapid design of articulated objects. *ACM Trans. Graph* 41, 4, Article 89 (2022), 8 pages.
- [10] Joon Hyub Lee, Hyunsik Oh, Junwoo Yoon, Seung-Jun Lee, Taegyu Jin, Jemin Hwangbo, and Seok-Hyung Bae. 2024. RobotSketch: an interactive showcase of superfast design of legged robots. In *ACM SIGGRAPH 2024 Emerging Technologies*. Article 17, 2 pages.
- [11] Tony Lewin and Ryan Borroff. 2010. *How to Design Cars Like a Pro*. Motorbooks.
- [12] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3D: high-resolution text-to-3D content creation. In *Proc. ICCV '23*. 3836–3847.
- [13] Stuart Macey and Geoff Wardle. 2014. *H-Point: The Fundamentals of Car Design & Packaging*. Design Studio Press.
- [14] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2022. SDEdit: guided image synthesis and editing with stochastic differential equations. arXiv:2108.01073 [cs.CV] <https://arxiv.org/abs/2108.01073>
- [15] Moreno Attilio Piccolotto. 1998. *Sketchpad+: Architectural Modeling through Perspective Sketching on a Pen-Based Display*. Master's thesis. Cornell University.
- [16] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2022. DreamFusion: text-to-3D using 2D diffusion. arXiv:2209.14988 [cs.CV] <https://arxiv.org/abs/2209.14988>
- [17] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with CLIP latents. arXiv:2204.06125 [cs.CV] <https://arxiv.org/abs/2204.06125>
- [18] Scott Robertson and Thomas Bertling. 2013. *How to Draw: Drawing and Sketching Objects and Environments from Your Imagination*. Design Studio Press.
- [19] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. arXiv:2112.10752 [cs.CV] <https://arxiv.org/abs/2112.10752>
- [20] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. 2022. Photo-realistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, Vol. 35. 36479–36494.
- [21] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proc. ICML '15*, Vol. 37. 2256–2265.
- [22] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2022. Denoising diffusion implicit models. arXiv:2010.02502 [cs.LG] <https://arxiv.org/abs/2010.02502>
- [23] Vizcom. 2023. *Vizcom*.
- [24] Qiang Wang, Di Kong, Fengyin Lin, and Yonggang Qi. 2023. DiffSketching: sketch control image synthesis with diffusion models. arXiv:2305.18812 [cs.CV] <https://arxiv.org/abs/2305.18812>
- [25] Min Xin, Ehud Sharlin, and Mario Costa Sousa. 2008. Napkin sketch: handheld mixed reality 3D sketching. In *Proc. VRST '08*. 223–226.
- [26] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. In *Proc. ICCV '23*. 3836–3847.