3D Sketching + 2D Generative AI for Car Exterior Design

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

Jeongche Yoon
Department of Industrial Design
KAIST
Daejeon, Republic of Korea
jeongche.yoon@kaist.ac.kr

Sang-Hyun Lee
Department of Industrial Design
KAIST
Daejeon, Republic of Korea
sang-hyun.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

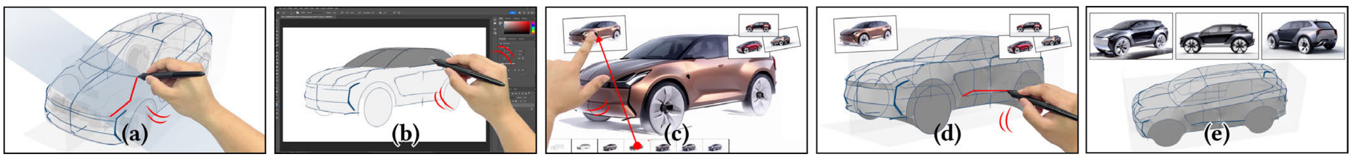

Figure 1: In a user study, a professional car designer used our novel interactive system to: (a) sketch a new exterior concept in 3D; (b) capture a hidden-line-removed image of the 3D sketch from a desired viewpoint, which was automatically opened in Photoshop for adding details and shading; (c) input the drawing into a 2D generative AI model to explore diverse design variations and collect inspiring images; (d) sketch new ideas directly onto the reconstructed 3D mesh using those images as reference; and (e) repeat this process to finally produce high-quality, view-consistent renderings from multiple viewpoints.

Abstract

Recent advances in 2D generative AI are finding applications in highly specialized fields, such as car exterior design. However, current 2D-centric approaches have several limitations: each viewpoint requires a new sketch; consistency across different viewpoints is difficult to maintain; and steering design development in the desired direction is challenging. To address these limitations, we propose a novel 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. In a formal user study, professional car designers used our system to create new exterior concepts across all major car body types. They produced high-quality, view-consistent sets of 2D renderings in a short time, demonstrating diverse patterns of progressive design development.

CCS Concepts

• **Human-centered computing** → **Interaction techniques**.

Keywords

3D sketching; 2D generative AI; car design

ACM ISBN 979-8-4007-2037-6/25/09
https://doi.org/10.1145/3746059.3747609

ACM Reference Format:
Seung-Jun Lee, Jeongche Yoon, Sang-Hyun Lee, Joon Hyub Lee, and Seok-Hyung Bae. 2025. 3D Sketching + 2D Generative AI for Car Exterior Design. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25), September 28–October 01, 2025, Busan, Republic of Korea.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3746059.3747609

1 INTRODUCTION

Thanks to rapid advances in diffusion-based 2D generative AI, car designers can now easily obtain realistic renderings from minimal sketches. This capability can help designers quickly explore a broad range of possibilities and develop them into concrete concepts during the early stages of design.

However, there are several limitations to applying 2D generative AI in professional car design practice. Once designers obtain a satisfactory rendering from one viewpoint, they must create a new sketch all over again to visualize the design from another angle, which requires additional time and effort. In addition, renderings generated from hand-drawn sketches at different viewpoints often lack view-consistency [15, 30, 65], resulting in a set of renderings that may fail to describe the same 3D shape. Furthermore, as designers repeatedly cycle through 2D sketching and 2D generation, design intent may be lost in translation [38], making it difficult to build on intermediate results and steer design development in the desired direction.

We suggest leveraging 3D sketching to help overcome these limitations. 3D sketching enables designers to intuitively express their ideas as 3D curves through minimal gestures, similar to those used in 2D sketching [5]. Once created, a 3D sketch can be viewed as a 2D line drawing from any desired viewpoint, and new ideas can be progressively added onto it. Given these strengths, 3D sketching has the potential to serve as a structural backbone in the 2D generative AI-based design process.

In this paper, we propose a novel workflow for car exterior design that integrates the strengths of 3D sketching and 2D generative AI in a complementary manner. This workflow enables designers to seamlessly transition between expressive 3D sketching, detailed 2D drawing, and realistic 2D generation, ensuring view-consistent and progressive development of designs across multiple viewpoints (Figure 1). We conducted a formal user study with professional car designers who used our system to design new exterior concepts across all major car body types, and demonstrated that our workflow can effectively support early-stage car design practice.

2 RELATED WORK

In this section, we first highlight the importance of sketching in early-stage car design. We then outline recent advances in diffusion-based 2D generative models and discuss their integration into design workflows, as well as their limitations in car design. Next, we introduce 3D sketching as a complementary method to guide 2D generative AI, and highlight techniques suitable for car exterior design. Finally, we review 3D mesh reconstruction methods needed for more intuitive and iterative 3D sketching.

2.1 Car Design Process

In the early design stage for a new car, designers define the car's target market segment, functional goals, and packaging, including its size, proportions, and layout [40]. Based on these specifications, they engage in intensive and competitive 2D sketching [37]. They rapidly conceive numerous ideas through concise line drawings, then create realistic renderings of selected concepts to communicate with colleagues [16]. Typically, designers create rendering sets from multiple viewpoints, most commonly front three-quarter, side, and rear three-quarter views, to describe the car's complex 3D form [9].

Through iterative design reviews, refined concepts are developed into scale or full-size clay models, and subsequently into sophisticated engineering CAD models for production [40]. Throughout this process, designers' 2D rendering sets serve as essential means for establishing design directions. This study aims to support the early concept design stage, where designers must rapidly explore a wide range of possibilities and visualize their ideas as view-consistent sets of 2D renderings within a short timeframe.

2.2 2D Generative AI

Recent advances in generative AI have opened new possibilities for accelerating car design processes. Particularly, diffusion models have shown remarkable success in image generation, and are now increasingly adopted in car design for initial concept exploration and rapid visualization.

Diffusion models generate images through an iterative process of noise addition and removal [19]. Researchers have improved and optimized this approach [53, 58], enabling the generation of high-quality results at remarkable speeds. These advancements have led to the emergence of models capable of generating realistic high-resolution 2D images of diverse styles from text prompts [44, 50, 55] and from 2D sketches [41, 63, 70].

Some studies have explored utilizing 2D generative AI for product design, including cars [11, 38]. However, these studies mainly focus on obtaining single-view renderings, having limited applicability in design contexts where sets of view-consistent 2D renderings from multiple viewpoints are required, such as car design.

A 3D reference can help ensure view-consistency when captured from multiple viewpoints and used as inputs for 2D generation. Several studies and commercial services have explored different approaches to creating such references in the context of 2D generative AI-based design, but with their own shortcomings.

Physical prototyping [68, 69], such as foam cutting or paper modeling, suits simple, everyday products but not fluid, complex forms of car exteriors. 3D CAD modeling [39, 45, 61] can produce such intricate forms, but it may cause a bottleneck during rapid idea exploration in the early design stages [10]. Image-to-3D mesh generation [62] is simple and fast, but it produces incompatible meshes when given 2D images of the same car viewed from different viewpoints, since each mesh fails to incorporate visual features that are not visible from its specific viewpoint. Instead, this study lets designers leverage their most powerful weapon, their drawing skills, to produce high-quality, presentation-ready renderings of creative concepts from compelling angles, as quickly and easily as possible.

2.3 3D Sketching

3D sketching is an interaction 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 [5]. The resulting 3D sketch can not only be viewed from any angle but also reviewed at real scale in VR [33], and used as reference data to accelerate subsequent modeling and engineering stages [5, 28, 29, 31].

Thanks to its ease of creation and richness in spatial information, 3D sketching can guide and complement 2D generative AI in design workflows. Prior work [56] demonstrated this potential in 3D scene design, where users can arrange 3D cards with scene objects sketched on them and use AI to generate images of the whole scene. In contrast, our previous work [36] focused on professional car design practice, enabling designers to express complex car forms using a set of 3D sketching methods. In this study, we polished our system and added on-mesh 3D sketching, as well as conducting a formal user study with multiple car designers.

The most commonly used 3D sketching method involves projecting 2D curves onto predefined 3D planes, resulting in 3D planar curves [17, 23, 28, 29, 31, 49, 64]. This approach aligns well with the typical practice of car designers, who often use orthographic 3D planes to establish overall proportions from canonical views such as front, side, rear, and top [17]. While 3D planar curves are sufficient for capturing basic forms, 3D spatial curves are required to represent the sophisticated and nuanced forms of cars [16]. Thus, our system incorporates additional 3D sketching methods [5], including the orthographically extruded surface method, the single-view symmetric epipolar method, and the two-view epipolar method for sketching 3D spatial curves, alongside the orthographic plane method for sketching 3D planar curves.

Finally, our comprehensive set of 3D sketching methods includes using 3D meshes as projection surfaces [3, 18], naturally extending a common practice of car designers who frequently use pictures of existing cars as underlays to effortlessly draw on top of and reimagine them in new ways. This method works not only with readily available 3D meshes of existing cars but also with 3D meshes rapidly and faithfully reconstructed from 3D sketch curves previously drawn by designers themselves. Reconstructing 3D meshes from 3D sketch curves, creating additional sketch curves using the meshes, and repeatedly updating the meshes based on newly created sketch curves can lead to a virtuous cycle of progressive design development.

2.4 3D Mesh Reconstruction

A common approach to reconstructing 3D meshes from 3D curves involves detecting closed loops within 3D curve networks and surfacing these loops [1, 2, 8, 46, 47, 54, 60, 71]. However, designers' sketches are often not tightly connected and may lack sufficient closed loops. Although techniques have been proposed to form well-connected curve networks by applying geometric corrections to the underlying 3D curves [67], such corrections may unintentionally distort designers' original intent.

Consequently, studies have explored methods to reconstruct 3D meshes from 3D sketches without relying on fully connected curve networks. One approach converts 3D curves into a voxel grid, generates a gradient vector field from the grid, and converges an initial spherical mesh onto the 3D curves [4]. However, due to resolution reduction during voxelization, this method struggles to accurately reconstruct complex surfaces. Another approach segments a coarse initial mesh according to 3D curves, then applies surface fitting to each segment, distinguishing smooth and sharp edges [66]. However, this method is computationally intensive and often requires several minutes of calculation for high-resolution sketches, hindering real-time design iteration.

Yet another approach reconstructs surfaces from point clouds sampled along 3D curves. Unlike classical point cloud reconstruction methods [26, 27] that assume dense and uniform point distributions, such as those obtained from 3D scans, recent approaches can effectively reconstruct smooth surfaces even from sparse and non-uniform point clouds, such as those sampled from sketch curves [20, 21, 48]. In particular, iterative Poisson surface reconstruction (iPSR) [20] achieves high processing speed by initially assigning random normals, performing Poisson reconstruction, and then iteratively refining normals based on the preliminary surfaces generated. We utilize iPSR to automatically reconstruct high-fidelity meshes from 3D sketch curves within a few tens of seconds.

3 SYSTEM

We propose a novel interactive system that seamlessly integrates 3D sketching, 2D drawing, and 2D generation for the early stage of car exterior design (Figure 2). Our system utilizes 3D sketches as structural backbones to ensure view-consistency among generated 2D renderings and to enable designers to progressively steer AI-based car exterior design in the desired direction.

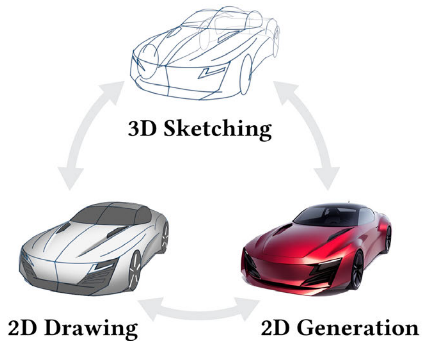

Figure 2: Within our workflow, a car designer can seamlessly transition between expressive 3D sketching, detailed 2D drawing, and realistic 2D generation for view-consistent and progressive design development.

3.1 Workspaces

Our proposed design workflow spans multiple dimensionalities (2D and 3D) and viewpoints. To facilitate seamless transitions between them, our system integrates three specialized types of workspaces: one 3D workspace for 3D sketching, multiple 2D workspaces for 2D drawing and 2D generation, and one on-screen workspace for managing visual references (Figure 3).

3D workspace. In the 3D workspace, designers can navigate freely using multi-touch gestures and perform 3D sketching on top of a car body template (Figure 3a) comprising a *3D underlay* and a *3D grid*. The 3D underlay is a semi-transparent monochrome 3D mesh upon which designers can directly sketch 3D curves, analogous to how car designers use images of existing cars as underlays to sketch new designs upon established proportions [52]. The 3D grid comprises four orthographic planes (front, rear, center, and ground) [17] aligned to the 3D underlay, and designers can use the planes to construct custom sketch planes and surfaces for 3D sketching.

2D workspace. Designers can capture 3D sketches from desired viewpoints, thereby creating new 2D workspaces (Figure 3b), and perform 2D drawing and 2D generation on top of them (Figure 3c). Because 3D sketches often contain many overlapping curves, the system automatically applies hidden-line removal upon creation of a 2D workspace, yielding clean and simplified line drawings (Figure 3b). Furthermore, the system automatically organizes the images produced within each 2D workspace chronologically as *thumbnails* (Figure 3c), thus relieving designers from manual file management. 2D workspaces are visually represented as *flagstick* widgets within the 3D workspace, allowing designers to quickly revisit them as needed via a tap on a flagstick widget (Figure 3d).

On-screen workspace. Designers can use the entire screen as an interactive pin board analogous to those on studio walls to which car designers pin visual materials for reference. Designers can create *reference image* widgets that are overlaid onto 3D and 2D workspaces and easily arrange or resize them using multi-touch gestures (Figure 3c, d). This setup can help designers make informed design decisions based on visual references regardless of whether they are performing 3D sketching, 2D drawing, or 2D generation. Additionally, these widgets function as shortcuts to instantly navigate back to the associated 2D workspaces.

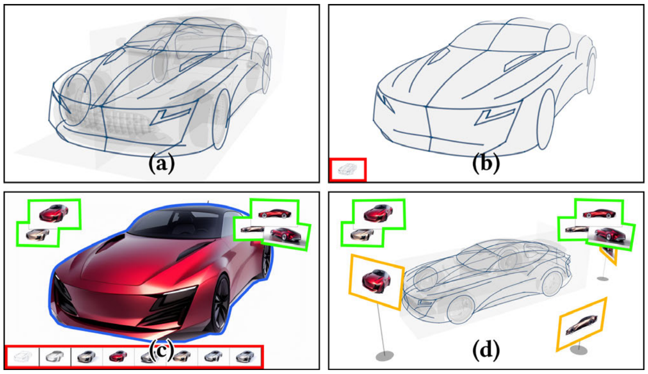

Figure 3: Workspaces. The designer can (a) perform 3D sketching within the 3D workspace and (b) create a 2D workspace from a desired viewpoint, which provides a clean 2D line drawing obtained through hidden-line removal applied to the 3D sketch. The designer's 2D drawing can serve as input for generating multiple renderings through 2D generative AI. (c) The 2D workspace shows all 2D outputs as *thumbnails* (outlined in red), which designers can tap to make fullscreen or drag onto the screen to create *reference images* (outlined in green). (d) When the designer rotates the view to return to the 3D workspace, *flagsticks* (outlined in orange) appear, marking the positions and orientations of the 2D workspaces. Tapping on a reference image or a flagstick instantly reopens the corresponding 2D workspace.

3.2 3D Sketching

Designers can utilize our minimal yet versatile set of 3D sketching methods to create essential 3D curves expressing car exterior designs (Figures 4-7). Designers can erase these curves partially or entirely, and the system automatically mirrors all curves, considering that most cars have symmetric 3D forms. In addition, the system highlights intersections between active sketch planes/surfaces and existing 3D underlays/curves with colored lines or dots to provide additional depth cues [5, 6].

3D sketching on a plane. There are various ways to define a 3D sketch plane, such as marking points that the plane should pass through [6, 29], extruding linear strokes drawn on canonical planes [64], or moving and rotating existing planes [17, 29, 31]. Car designers often prefer orthographic views (front, side, and top) for accurate proportions. Reflecting this practice, prior work [17] proposed an interface that allows users to move orthographic grid planes along the X, Y, and Z axes and sketch on them. Building on this approach, our system extends the method by incorporating touch gestures (Figure 4).

3D sketching on a surface. A 3D sketch surface is typically defined by extruding an existing curve in a specific direction. Prior work [5] proposed determining this direction using widgets and pen gestures. Our system builds on this approach by allowing designers to directly use existing grid planes for extrusion (Figure 5).

3D sketching on a mesh. Sketching directly on a mesh [3, 18] is a highly intuitive approach, yet it has seen limited use in design practice due to the difficulty of drawing clean lines on uneven surfaces. In our system, even if the initial underlay surface is complex and produces rough curves, designers can easily refine them later using the other sketching methods. Moreover, once a sufficient number of 3D curves have accumulated to express the intended 3D form, the underlay can be replaced with a watertight 3D mesh reconstructed from these curves, after which designers can effortlessly add well-defined 3D curves directly on this new surface (Figure 6).

3D sketching with two curves. Unlike the other three types of 3D sketching methods, which require explicit sketch planes, surfaces, or meshes to project curves onto, this approach allows designers to draw a curve twice, freely in the air, to uniquely situate a 3D spatial curve. The two curves may represent mirrored curves drawn from a single viewpoint [7], or the same curve drawn from two distinct viewpoints [24, 25]. Prior work [5] combined these two methods into a single, natural sequence of sketching interaction, which we adopt in our system (Figure 7). Our system further extends this interaction by allowing designers to draw multiple second curves, enabling them to iteratively create pairs of 3D spatial curves until they find one that best matches their intent.

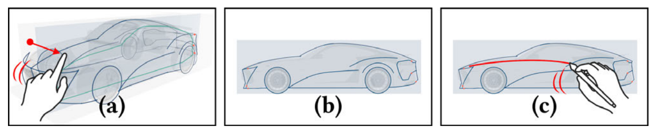

Figure 4: Orthographic plane method. The designer can (a) drag a grid plane (in this case, the center plane) to a desired position to activate a sketch plane, (b) double-tap it to enter an orthographic view (in this case, the side view), and (c) sketch 3D planar curves onto the sketch plane.

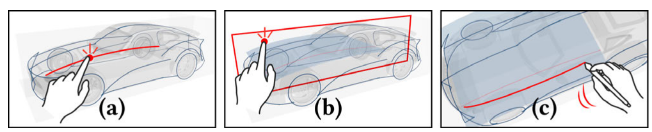

Figure 5: Orthographically extruded surface method. The designer can (a) tap one or more 3D curves to use as a profile, (b) tap a grid plane (in this case, the center plane) whose normal aligns with the desired extrusion direction to activate a sketch surface, and (c) sketch 3D spatial curves onto it.

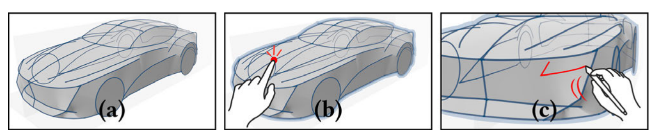

Figure 6: Mesh projection method. The designer can (a) press a keyboard shortcut to replace the underlay with a reconstructed 3D mesh, (b) tap the underlay to activate projection, and (c) sketch 3D spatial curves onto it.

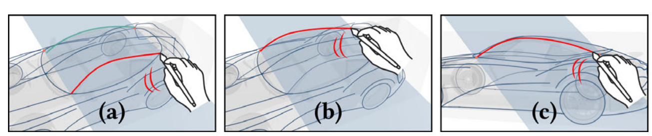

Figure 7: Two-curve method. The designer can (a) draw the first curve, then either (b) draw a symmetric counterpart from the same viewpoint or (c) redraw the curve from another viewpoint to create 3D spatial curves.

3.3 3D Mesh Reconstruction

In our system, reconstructed 3D meshes serve two roles: first, as iteratively updated 3D underlays for 3D sketching, and second, as temporary geometries for hidden-line removal. Since they have different requirements for mesh quality and reconstruction speed, we employ two different algorithms accordingly.

3D underlay for 3D sketching. 3D underlays must precisely fit designers' sketch curves and be watertight to function as sketch surfaces. The reconstruction process consists of three steps (Figure 8). First, we sample 3D points from the polyline-based sketch curves to generate a point cloud. Second, we reconstruct a mesh using the iPSR algorithm [20], setting the octree depth to 9 and the neighborhood parameter to 100, based on empirical observations. Third, we apply post-processing to achieve a tighter fit, which includes increasing mesh resolution through subdivision, fitting mesh vertices near sketch curves onto the nearest points on those curves, and applying Laplacian smoothing [59] to the fitted surface. The entire procedure takes approximately 30 seconds in our implementation.

Hidden-line removal. When using a 3D sketch of a car as an underlay for 2D drawing, showing all curves can be confusing, as some should be hidden by occlusion. However, since there is no volume to determine visibility, we generate a custom mesh to approximate the car's shape and remove hidden curves in real time.

This mesh does not need the fidelity of the aforementioned 3D underlay, but it must be quick to reconstruct and still accurately capture the overall silhouette of the 3D sketch. The reconstruction process based on heuristics regarding car body shapes consists of three steps (Figure 9). First, for the relatively cylindrical front (ahead of the front-wheel axis) and rear (behind the rear-wheel axis) sections, we sample cross-sections radially and connect them into triangular strips. Second, for the relatively rectangular central section, we sample cross-sections linearly along the car's length and connect them into parallel strips. Third, the three meshes are merged. The entire procedure takes a few hundred milliseconds in our implementation.

3.4 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, is expressed more efficiently (Figure 10). Professional car designers develop deep familiarity and expertise with specialized 2D image editing tools through years of practice. Therefore, instead of duplicating these tools, we provide seamless integration, so designers can easily export images, modify them in their preferred software, and automatically re-import them.

3.5 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 11). These high-quality 2D renderings can then serve as a basis for further 3D sketching, 2D drawing, and 2D generation, facilitating progressive and 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 and preferences.

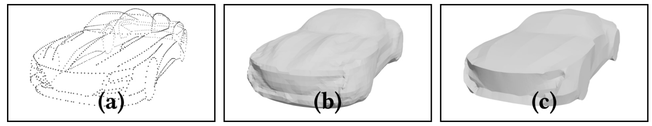

Figure 8: 3D mesh reconstruction for 3D sketching. The system (a) extracts a point cloud from the 3D sketch, (b) generates an initial mesh using the iterative Poisson surface reconstruction (iPSR) algorithm [20], and (c) applies post-processing.

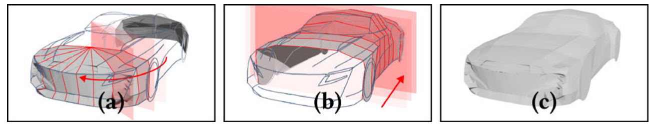

Figure 9: 3D mesh reconstruction for hidden-line removal. The system (a) patches cylindrical section curves for the front and rear, (b) patches parallel section curves for the middle, and (c) merges three resulting meshes into one.

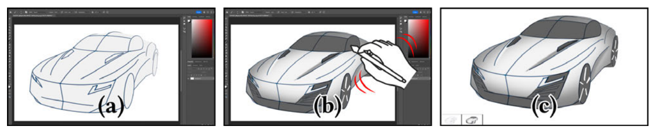

Figure 10: 2D drawing. When the designer presses a keyboard shortcut within the 2D workspace, (a) the system automatically launches an external 2D editing tool (in this case, Photoshop) and loads the current image, (b) allowing the designer to proceed with 2D drawing. (c) Upon returning to the system, the designer can find the drawing automatically updated in the 2D workspace.

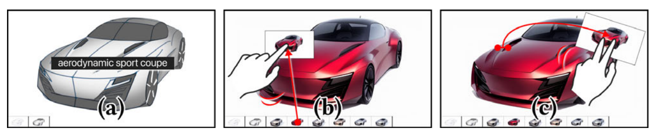

Figure 11: 2D generation. (a) The designer can configure parameters including text prompts, drawing influence levels, and the number of 2D images to generate. (b) After the 2D generation is completed, resulting in a row of thumbnails, the designer can drag-and-drop a thumbnail to create a reference image and (c) freely move, rotate, and resize it.

4 IMPLEMENTATION

We implemented our system using the Unity 3D game engine. Although our workflow is conceptually compatible with any 2D drawing tool or generative model, we chose state-of-the-art applications widely adopted in professional design practice. 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 [62], a 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 RTX 4070 GPU, paired with a Wacom Cintiq Pro 24 Touch digital tablet supporting both multi-touch and pen inputs.

5 USER STUDY

We conducted a formal user study to evaluate the benefits of our proposed system for car exterior design practice. We invited six professional designers with extensive experience in car design to participate in a workshop, where they were tasked with creating new design concepts covering all twelve standard car body types [13], ranging from superminis to pick-ups. The designers produced high-quality results using our system, while we gathered detailed data on their workflows and experiences.

5.1 Participant

Since our aim was to investigate whether our system could provide practical benefits to professional car designers, we only recruited designers who were working in the car industry to participate in our in-depth, hands-on workshops.

The participants were six professional car exterior designers, each with work experience averaging 7.5 years (min: 3 years, max: 14 years) in the industry. At the time of our study, they were working as junior (P2, P4, P5) or senior (P1, P3, P6) designers at international car design studios (P1, P2, P3) or car manufacturers (P4, P5, P6). Their expertise covered various car body types, including superminis (P1, P2, P3, P4, P5, P6), small off-roads (P3, P6), small MPVs (P1, P3), small family cars (P4, P6), pick-ups (P6), large off-roads (P1, P2, P3), large MPVs (P3, P4, P5), and business and family vans (P3), having led or participated in an average of 9.5 concept and production projects. All but one (P3) had experience using 2D generative AI in their work.

5.2 Task

The designers were asked to use our system to create new design concepts for two car body types randomly assigned from the twelve defined by the European New Car Assessment Programme (Euro NCAP) [13] and produce a set of 2D renderings for three key views (front three-quarter, side, and rear three-quarter). They had up to two hours per car, which was determined as sufficient time through internal pilot testing. For each task, the designers were provided with an existing car's 3D model that they requested as the initial template (Table 1).

5.3 Procedure

The workshop comprised three sequential sessions. First, in a one-hour tutorial session, the experimenter followed a script to guide the designers through all system features step-by-step. Second, in a four-hour task session, the designers completed their design tasks, and were provided with a single-page cheat sheet covering all system features. Finally, in an individual one-hour debriefing session, each designer filled out a survey and revisited the design process with the experimenter by reviewing time-lapses of the recorded videos together and discussing their thoughts and results at each step of the process.

5.4 Measurement

We measured the frequency, duration, and order of each activity pertaining to 3D sketching, 2D drawing, and 2D generation. Specifically, 3D sketching comprised navigating, sketching, and waiting for mesh generation to complete; 2D drawing comprised adjusting viewpoints and drawing in Photoshop; and 2D generation comprised configuring parameters, waiting for generation to complete, reviewing generated images, and managing reference images (Figure 13). 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. Satisfaction with the system's key features and overall user experience were assessed through a 5-point Likert scale survey.

5.5 Result

Six designers used our system for 16 hours in total and successfully created new exterior design concepts for all twelve car body types, resulting in 12 final 3D sketches and 36 final 2D renderings (Figure 12), spending 1 hour and 20 minutes per car on average. They spent 44.8% of time on 3D sketching, producing 2.8 3D sketches per car. They spent 37.1% of time on 2D drawing, producing 13.3 2D drawings per car. Finally, they spent 18.1% of time on 2D generation, producing 46.3 2D renderings per car (Table 1). Survey responses indicated high satisfaction across usability with an average score of 4.5 (±2 SE: 0.1) on Q1-20, and across usefulness with an average score of 4.5 (±0.2) on Q21-25 (Figure 14).

Participant	Body type	Template	Time spent (h:mm (%))				Output (#)			3D curve (# (%))				
			3DS	2DD	2DG	Total	3DS	2DD	2DG	OPM	OESM	MPM	TCM	Total
P1	Pick-up	Ford F-150	00:29 (40.1)	00:28 (38.2)	00:15 (21.7)	01:12 (100)	2	20	76	278 (65.3)	66 (15.5)	50 (11.7)	32 (7.5)	426 (100)
	Large MPV	Volkswagen ID. Buzz	00:23 (34.2)	00:27 (41.0)	00:16 (24.8)	01:06 (100)	1	23	81	125 (50.8)	30 (12.2)	18 (7.3)	73 (29.7)	246 (100)
P2	Large off-road	Mitsubishi Outlander	00:58 (50.2)	00:50 (43.4)	00:07 (6.4)	01:55 (100)	4	12	28	22 (14.5)	0 (0.0)	102 (67.1)	28 (18.4)	152 (100)
	Roadster	BMW M4	00:21 (41.0)	00:19 (36.4)	00:11 (22.6)	00:51 (100)	4	9	48	56 (11.0)	4 (0.8)	448 (88.2)	0 (0.0)	508 (100)
P3	Large family car	Audi A5	00:40 (52.8)	00:19 (24.6)	00:17 (22.6)	01:16 (100)	7	16	44	55 (12.4)	291 (65.8)	96 (21.7)	0 (0.0)	442 (100)
	Supermini	Hyundai i20	00:41 (51.5)	00:13 (16.4)	00:26 (32.1)	01:20 (100)	4	14	84	136 (27.8)	274 (55.9)	66 (13.5)	14 (2.8)	490 (100)
P4	Business and family van	Benz V Class	00:55 (66.5)	00:23 (27.7)	00:04 (5.9)	01:22 (100)	3	11	19	69 (45.1)	22 (14.4)	62 (40.5)	0 (0.0)	153 (100)
	Sport	Porsche 911	00:40 (45.8)	00:34 (37.6)	00:14 (16.6)	01:28 (100)	1	16	46	77 (29.7)	80 (30.9)	92 (35.5)	10 (3.9)	259 (100)
P5	Small off-road	Kia Sportage	00:33 (36.2)	00:44 (47.8)	00:15 (16.0)	01:32 (100)	2	14	28	91 (41.6)	0 (0.0)	128 (58.4)	0 (0.0)	219 (100)
	Small family car	Volkswagen Golf	00:18 (28.3)	00:40 (57.8)	00:09 (13.9)	01:07 (100)	1	10	26	220 (49.1)	0 (0.0)	228 (50.9)	0 (0.0)	448 (100)
P6	Executive	Tesla Model S	00:31 (38.9)	00:32 (39.7)	00:17 (21.4)	01:20 (100)	3	8	43	69 (27.9)	16 (6.5)	160 (64.8)	2 (0.8)	247 (100)
	Small MPV	Fiat 500X	00:48 (52.2)	00:32 (34.4)	00:12 (13.4)	01:32 (100)	2	6	33	71 (28.1)	48 (19.0)	132 (52.2)	2 (0.8)	253 (100)
Average			00:36 (44.8)	00:30 (37.1)	00:14 (18.1)	01:20 (100)	2.8	13.3	46.3	105.8 (33.6)	69.2 (18.4)	131.8 (42.6)	13.4 (5.4)	320.2 (100)

Table 1: Body type, initial template, time spent (3DS: 3D sketching, 2DD: 2D drawing, 2DG: 2D generation), number of outputs, and number of 3D curves created by method (OPM: orthographic plane method, OESM: orthographically extruded surface method, MPM: mesh projection method, TCM: two-curve method) for each task, ordered by participant.

Figure 12: The 3D sketches and 2D renderings produced by the designers in our user study, from the front three-quarter, side, and rear three-quarter views of (a) a pick-up by P1, (b) a large MPV by P1, (c) a large off-road by P2, (d) a roadster by P2, (e) a large family car by P3, (f) a supermini by P3, (g) a business and family van by P4, (h) a sport by P4, (i) a small off-road by P5, (j) a small family car by P5, (k) an executive by P6, and (l) a small MPV by P6.

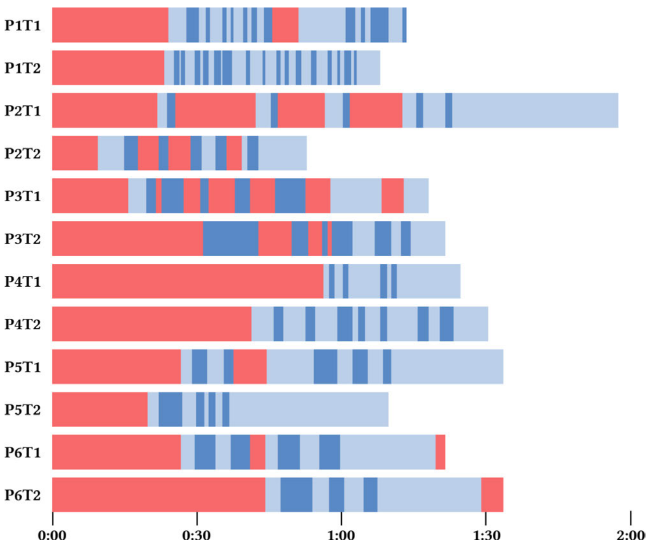

Figure 13: Timeline of the designers' activities for all 12 design tasks. Each bar represents a design activity (red: 3D sketching, light blue: 2D drawing, dark blue: 2D generation), with the length proportional to the time spent. The designers demonstrated various patterns of progression and iteration.

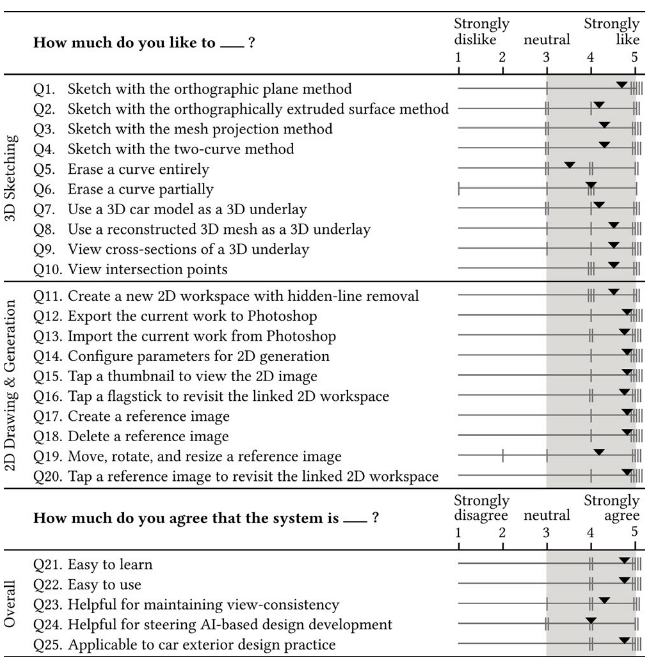

Figure 14: The 5-point Likert scale survey on the system's usability (Q1-20) and usefulness (Q21-25). Each tick corresponds to a designer's response, and each triangle to the average score. An average score of 4.5 (±2 SE: 0.1) across all questions indicates a satisfying overall experience.

6 DISCUSSION

We discuss how the designers utilized 3D sketching, 2D drawing, and 2D generation in a complementary manner to achieve satisfactory outcomes, highlighting the benefits of our integrated and flexible workflow. First, we examine how the four 3D sketching methods supported diverse modes of idea expression. Then, we explain how the designers achieved view-consistency through progressive development. Finally, we reflect on their experience collaborating with AI in our integrated system and consider its applicability to professional car design practice.

6.1 3D Sketching

The designers spent 44.8% of their time on 3D sketching, producing an average of 2.8 3D sketches per car, with final 3D sketches containing an average of 320.2 3D curves (Table 1). They reported high satisfaction with 3D sketching in our system and agreed that these methods allowed them to fluently express the intended 3D shapes of the cars, giving an average score of 4.3 (±0.2) on Q1-10 (Figure 14). We examine how each method was used during the tasks.

Mesh projection method (MPM). The most frequently used method was MPM (42.6%, Table 1), serving two main purposes. First, it was used in the early stages for rapid and freeform shape exploration on top of the 3D underlay as if *"doodling on paper"* (P4, P6). One designer described it as an *"ideal sketching method for exploring form when the initial idea hasn't been fully developed yet"* (P4). Second, it was used to add detailed graphic elements such as lamps, grilles, door lines, and character lines onto reconstructed 3D meshes once the overall form had been established. One designer noted, *"it was effective for sketching garnishes that attach to the main volume of the car"* (P3).

Orthographic plane method (OPM). The second most frequently used method was OPM (33.6%, Table 1). OPM was mainly used to define key proportions and volumes of the cars, which the designers described as a combination of CAD precision and sketching intuition (P1, P3, P5). The designers also noted that OPM was particularly useful when *"defining overall proportions from the side view in the early stages"* (P1), *"sketching cross-sections to define major volumes of a car"* (P2, P4), and *"drawing 3D curves at precise locations"* (P1, P6).

Orthographically extruded surface method (OESM). The third most frequently used method was OESM (18.4%, Table 1), and it was mainly used for two purposes. First, it was used to develop existing 3D planar curves into 3D spatial curves. One designer emphasized, *"because the 3D spatial curves were developed from planar curves drawn with consideration for the overall proportions and volumes, the design intent was automatically carried over"* (P1). Second, it was used to transform curves drawn on meshes into 3D spatial curves extending beyond the mesh surfaces. One designer emphasized, *"it's difficult to draw 3D spatial curves from scratch, but by combining* [MPM and OESM]*, I was able to trace curves and gradually refine them"* (P6).

Two-curve method (TCM). The least frequently used method was TCM (5.4%, Table 1). The designers noted that the method *"takes some time to get used to"* (P1, P5, P6). At the same time, they acknowledged the unique advantages of TCM, noting that it is essential for sketching *"key character lines"* (P3) or *"graphic *elements"* (P2) that must appear in a specific shape from a specific viewpoint, and *"useful for experimenting with forms that completely deviate from the 3D underlay"* (P4).

6.2 View-Consistency

The designers were able to produce high-quality, view-consistent 2D rendering sets using our system (Figure 12). They agreed that it helped maintain consistency across multiple viewpoints, giving an average score of 4.3 (±0.7) on Q23 (Figure 14). They estimated that achieving similar outcomes using conventional workflows with only 2D drawing and generation would take 3-4 hours (P1, P3, P6), or even 5-6 hours (P2, P4, P5). We discuss how our system facilitated view-consistent design development.

Minimizes repetitive hand sketching. Although the task given was to create renderings from three specific viewpoints, the designers produced drawings from an average of 5.8 viewpoints per car. One designer explained, *"each time the design was updated, I wanted to find another view that best presented the new design"* (P6). The designers noted that even experienced car designers find it *"tedious"* (P1, P6), *"difficult"* (P2, P4, P5), and *"time-consuming"* (P3) to draw accurate perspective sketches with high view-consistency from scratch every time. In contrast, they highlighted the benefits of our workflow, noting that *"with a single 3D sketch, I could produce as many inputs for generative AI as I wanted, from any viewpoint"* (P1), *"it was very convenient not having to redraw everything"* (P3), and *"being able to switch to 2D at any time gave me confidence to explore my design from various viewpoints"* (P5).

Eliminates the manual work of erasing hidden lines. The designers reported high satisfaction with the automatic hidden-line removal performed when creating a 2D workspace, giving an average score of 4.5 (±0.4) on Q11 (Figure 14). One designer emphasized, *"it was so natural that I didn't even realize the feature existed"* (P5). Regarding the hidden-line-removed 3D sketches, the designers noted, *"just adding some details and shading was enough to use it as input for 2D generation"* (P3), *"when starting 2D drawing, the entire intended form was already there in Photoshop, which made sketching easy"* (P4), and *"you can get usable results even by directly using a hidden-line-removed 3D sketch without any additional touches"* (P5).

Streamlines current industry practices to achieve view-consistency. The designers explained that current industry practice for achieving view-consistency across renderings typically involves creating a 3D CAD model and capturing it from various viewpoints to use as underlays. Based on this experience, they highlighted the benefits of our system, noting, *"since 3D sketching is much lighter than CAD, I was able to quickly express a 3D form"* (P1), *"I could define the shape in the early stages with dramatically improved speed"* (P4), and *"I could quickly create an initial 3D volume that accurately reflected the intended proportions and volumes"* (P5).

6.3 Progressive Development

All designers utilized the full range of system features (Figure 15), but applied them through different strategies (Figure 13), resulting in three distinct patterns of progressive development: *single-rounder* (Figure 16), *multi-rounder* (Figure 17), and *continuous-switcher* (Figure 18). We identify the characteristics of each pattern, explore the motivations behind each strategy, describe how progressive design development unfolded in each case, and discuss how our system supported these diverse processes.

Single-rounder: essentials in 3D, details in 2D. This pattern is characterized by a clear separation between an initial 3D sketching phase and a subsequent phase focused on 2D drawing and generation (Figure 16). It was observed in 5 out of the 12 tasks (P1T2, P4T1, P4T2, P5T2, P6T2). The designers working in this pattern established the main form of the cars through 3D sketching, then progressively added details through 2D drawing and generation to produce the final rendering sets.

The designers emphasized that they approached these particular tasks with a relatively clear vision of the overall 3D form of the cars. By the end of the initial 3D sketching phase, they felt confident that the essential structure had been fully captured, noting, *"the fundamental structure I had in mind was fully expressed"* (P1), *"the overall silhouette and main graphics were all there"* (P4), and *"all the important curves were already defined"* (P5). This solid foundation

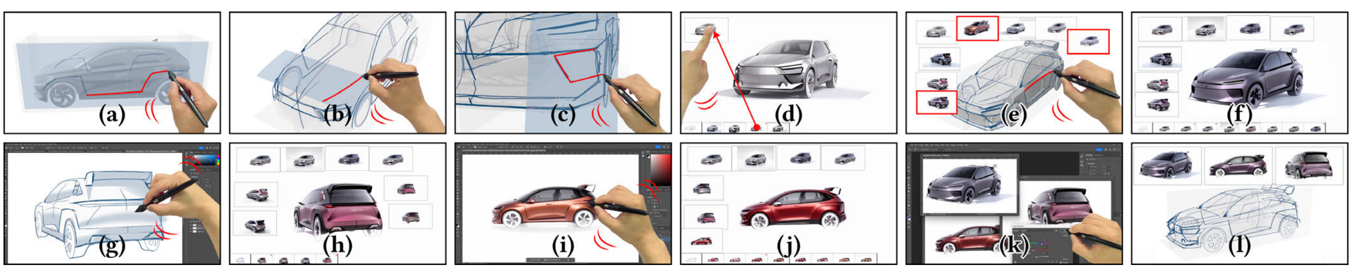

Figure 15: The designers successfully completed the given task using all system features. For example, P3 began his second task (P3T2) by sketching the overall shape using (a) the orthographic plane method, (b) the orthographically extruded surface method, and (c) the two-curve method. (d) P3 then explored design ideas by generating images from multiple viewpoints, and (e) after selecting preferred results (outlined in red), updated the 3D sketch on a reconstructed 3D mesh using the mesh projection method with reference to those images. (f) After finalizing the front three-quarter view, P3 used it as an anchor to further develop the other views. For the rear three-quarter view, (g) P3 added simple shading to a hidden-line-removed image of the 3D sketch and (h) generated a refined output. For the side view, (i) P3 iteratively generated variations, selected and refined the best result, and repeated the process until (j) achieving a satisfactory outcome. Finally, (k) P3 adjusted color and details across three viewpoints to produce (l) a high-quality, view-consistent rendering set.

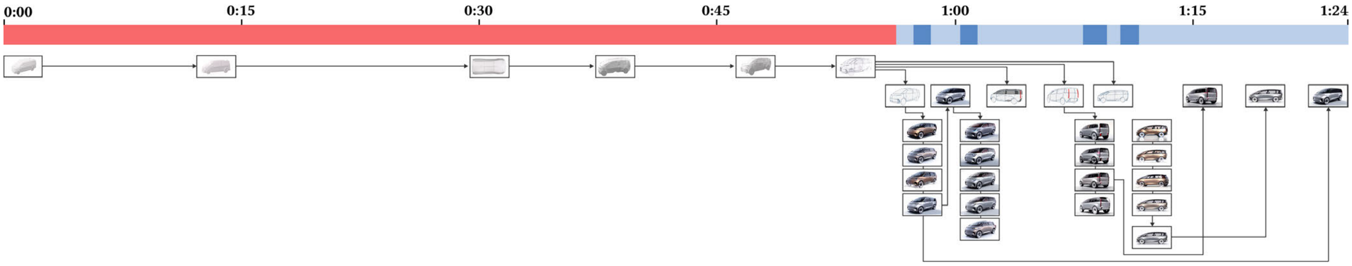

Figure 16: A task that showed progressive development pattern 1: *single-rounder* (P4T1). Each bar represents a design activity (red: 3D sketching, light blue: 2D drawing, dark blue: 2D generation). In the initial 3D sketching phase, P4 developed the overall shape of the car through iterative 3D sketching and 3D mesh reconstruction. In the later 2D phase, P4 added details and finalized the rendering set through iterative 2D drawing and 2D generation.

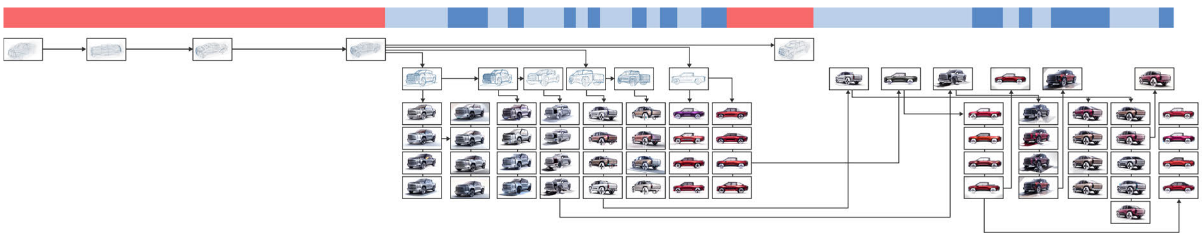

Figure 17: A task that showed progressive development pattern 2: *multi-rounder* (P1T1). P1 quickly built the overall shape through 3D sketching, then created renderings for all three views to review the design. As new ideas emerged, P1 refined the 3D sketch and produced the final rendering set again, based on the updated 3D sketch.

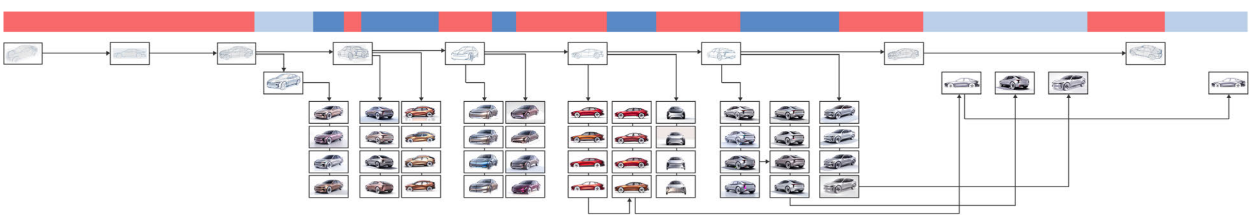

Figure 18: A task that showed progressive development pattern 3: *continuous-switcher* (P3T1). P3 developed the design in an exploratory way by obtaining new ideas from 2D generation at each viewpoint, integrating them into the 3D sketch, and repeating it multiple times until he collected desirable images for all viewpoints.

enabled them to focus on developing graphic details during the subsequent 2D phase.

In particular, the designers strategically leveraged 3D mesh reconstruction throughout the initial 3D sketching phase. They followed an iterative workflow by sketching 3D curves, generating a new 3D underlay from those curves, evaluating the resulting shape, and continuing to sketch on the updated surface. This process enabled progressive development of the overall form until they achieved their intended design. One designer emphasized, *"the 3D mesh faithfully reflected the 3D sketch and was generated quickly, so I could continuously update the underlay and refine the form"* (P6).

Multi-rounder: one pass, then refine. This pattern involves two or more rounds of the sequential 3D sketching-2D drawing-2D generation cycle (Figure 17), and was observed in 3 of the 12 tasks (P1T1, P5T1, P6T1). The designers working in this pattern first developed concrete concepts in the initial round. They then incorporated new ideas that emerged during that round into a second round of 3D sketching, followed by 2D drawing and generation, to complete the rendering sets for the second time, and so on.

The designers described their approach as follows. One designer noted, *"I first created a draft of the three-view rendering set to get a quick sense of the overall mood. I really liked the shape of the headlight and wheel arch in the side view rendering, so I updated the 3D sketch to match them"* (P1). Another designer explained, *"I tried to ground my initial design on the front three-quarter view, and then incorporated the details of the headlight and side character lines from that view into the 3D sketch"* (P5). Yet another designer reflected, *"I didn't even sketch that much in 3D at first, but I got a form that I liked, so I moved right into the 2D phase. When I saw that the graphics at the front of the car in the generated side view fit well with the overall shape, I updated the 3D sketch accordingly"* (P6).

In particular, the designers emphasized the value of quickly visualizing front and rear three-quarter views for arriving at a concrete concept early on, which made subsequent iterations possible. One designer noted, *"the overall impression of a car is determined by its front and rear views"* (P1), while another designer pointed out that *"it is common practice in 2D workflows to begin from the side view, rather than perspective views which can be distorted or inconsistent"* (P5). In contrast, the designers found that our system enabled them to *"quickly review realistic, view-consistent front and rear three-quarter views"* (P6), which helped them *"evaluate and determine the design direction in earlier phases"* (P5).

Continuous-switcher: 3D and 2D, back and forth. Unlike the other patterns where design development occurs in clearly divided sequential rounds, this pattern is characterized by frequent switching between 3D sketching, 2D drawing, and 2D generation, resulting in a more organic progression (Figure 18). This pattern was observed in 4 of the 12 tasks (P2T1, P2T2, P3T1, P3T2). The designers working in this pattern, jumped fluidly across multiple viewpoints after initial 3D sketching, alternating between 2D drawing and 2D generation while continuously feeding new ideas back into 3D sketches, culminating in the final 2D rendering sets.

The designers explained that they approached these tasks with open-ended starting points and embraced a more exploratory, iterative design process, noting they *"started with only an abstract form and mood"* (P2), *"defined only the lamp positions, hard points, and simple character lines before beginning to explore ideas"* (P2), *"had the desired silhouette and intended to develop the graphics in 2D"* (P3), and *"had the shape I wanted, but no specific concept, so I turned to AI for ideas"* (P3).

In particular, the designers frequently used hidden-line-removed 3D sketches as direct inputs for 2D generation, since they *"clearly conveyed the intended silhouettes"* (P3). They proceeded directly to 2D generation without an intermediate 2D drawing step in 20 out of 31 instances (64.5%), unlike the single-rounder and multi-rounder patterns, which consistently included a 2D drawing phase. Because the designers began with rough forms and actively sought new ideas from 2D generative AI, they required frequent transitions between 3D sketching and 2D generation, with hidden-line removal playing a critical role in enabling this fluid workflow.

Interestingly, similar patterns were reported in prior work that combined physical prototyping with 2D generative AI [69]. In both studies, designers varied in when they first engaged generative AI and how frequently they transitioned between 2D and 3D, depending on the clarity of their ideas and the level of refinement. Particularly, when designers lacked clear direction, as in our continuous-switcher pattern, they more readily turned to AI for help, resulting in more frequent switching between 2D and 3D. Such a parallel further highlights the importance of frictionless transitions across dimensionalities and tools to support diverse modes of creativity.

6.4 AI Collaboration

The designers exhibited different behavioral patterns across tasks, but a common theme emerged: while overall forms were developed manually through 3D sketching, details were refined with the help of 2D generative AI. We present the designers' experiences and explore the potential of this collaborative design process with AI.

They developed designs with a sense of authorship. The designers agreed that our system was helpful for steering AI-based design development, giving an average score of 4.0 (±0.7) on Q24 (Figure 14). They felt in control throughout the design process, noting, *"it felt like I was the design director, telling an AI designer what to do and receiving an unlimited stream of design results"* (P3) and *"I gave it a 3D sketch that described what I wanted, and it gave me outputs that matched my expectations"* (P4). These reflections indicate that, unlike conventional 2D generative-AI workflows often likened to using a "slot machine" [38, 51], our system enables designers to actively express and explore ideas through 3D sketching, then use the resulting 3D sketch as a structural backbone that guides and complements 2D generative AI.

AI-generated ideas stimulated design exploration. The designers noted that before a concrete form had been established, viewing *"a wide range of variations generated from minimal line sketches"* (P6) was *"useful for determining the overall tone and manner of the design"* (P4). As the process progressed and overall forms became more defined, AI-generated results supported the exploration of finer details within established structures. The designers highlighted several specific benefits, including *"different surface interpretations within the same silhouette"* (P6), *"character lines that simplified complex strokes"* (P1), *"unexpected but harmonious body colors"* (P1, P5), *"lamp designs I wouldn't have thought of"* (P2), and *"various combinations of garnishes that matched the intended style"* (P3). These examples demonstrate how AI-generated outputs served as valuable creative stimuli, supporting the designers in both early-stage ideation and later-stage refinement.

Designer's touch remains essential. The designers primarily used 2D drawing for three purposes. First, they added silhouettes, shading, or details before 2D generation, as these elements are difficult to express using 3D curves alone and are more efficiently handled in 2D (Figure 15g, h). Second, they refined promising images before generating more variations of them so that the noise is not amplified (Figure 15i, j). Third, they finalized the renderings by ensuring consistent color and details across all viewpoints (Figure 15k, l). Notably, this finalization step accounted for 17.3% of their total time, which amounted to approximately half of the total 2D drawing time. The designers emphasized the continued need for human intervention, noting, *"lines intended to express surfaces and reflections were sometimes rendered unnecessarily as parting lines or garnish details"* (P2), *"AI often failed to express bold designs because they were interpreted with traditional car design language"* (P3), and *"sometimes the output would feature random colors not matching the intended prompt"* (P5). These examples reaffirm the importance of human oversight and indicate that if generative AI continues to advance to the point where the need for refinement is minimized, the overall efficiency of our system could improve significantly.

6.5 System Integration

The designers transitioned between 3D sketching, 2D drawing, and 2D generation an average of 14.4 times per car (Figure 19), equating to approximately one transition every 5.6 minutes over a total duration of 80 minutes per car. We discuss how our system's tight integration benefited their design processes.

Seamless transitions supported deep focus. The designers reported becoming deeply immersed in their tasks while using our system, often attributing this immersion to the automated file handling, data transfer, and window management, giving an average score of 4.8 (±0.2) on Q12-16 (Figure 14). One designer noted, *"3D sketches were automatically transferred to Photoshop, and modified images in Photoshop were immediately sent to the 2D generation stage, so the workflow remained uninterrupted"* (P4). Another designer noted, *"In the past, I had to manually transfer data between* *different software and save files myself, but this automation allowed me to stay in the flow"* (P6). These comments highlight significant improvements in creative focus and workflow continuity.

The wall separating 2D and 3D was removed. A key benefit of seamless transitions was the elimination of rigid separations between 2D and 3D representations. The designers explained their traditional workflow challenges, noting, *"I usually first create a pseudo 3D model to use as a reference for 2D rendering. But, 3D CAD modeling tools are heavy, so once I move onto 2D, I don't feel like going back and fixing the 3D model even if I need to"* (P6), and *"I never even considered switching back and forth between 3D and 2D tools"* (P4). In contrast, when using our system, the designers noted, *"It felt natural to check the 2D generation result right after sketching in 3D, then jump back into 3D to make some quick fixes"* (P2), and *"I felt free to go back and forth between 2D and 3D at any time"* (P3). One designer characterized this transformed experience by describing our system as *"a bridge that connects 2D and 3D"* (P5).

Reference images supported design decisions. The designers collected an average of 12.3 reference images per car and reported high satisfaction with their use, giving an average score of 4.7 (±0.3) on Q17-20 (Figure 14). They particularly valued the intuitive and flexible arrangement of reference images, contrasting this with conventional workflows where they would typically *"print out reference images and pin them to the wall"* (P5) or *"open multiple windows on a secondary monitor"* (P3, P6). Moreover, the designers emphasized that reference images served not merely as visual cues but as a dashboard for monitoring overall design direction and progression. They noted, *"I felt reassured seeing the intermediate results accumulate step by step, confirming that my design was evolving in the direction I wanted"* (P4), *"I could review the history of key design decisions at a glance and immediately return to previous stages to reuse outcomes when necessary"* (P5), and *"with all intermediate and final results displayed in one view, I could continuously monitor the progression of my design"* (P6).

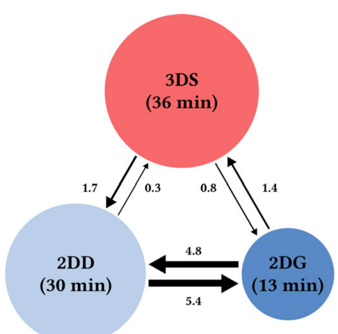

Figure 19: The designers transitioned between 3D sketching (3DS), 2D drawing (2DD), and 2D generation (2DG) an average of 14.4 times per car (min: 8, max: 33). The width of each arrow is proportional to the number of transitions, and the area of each circle is proportional to the time spent.

6.6 Applicability

The designers strongly agreed that our system is applicable to professional car design practice, giving an average score of 4.7 (±0.4) on Q25 (Figure 14). We discuss the main reasons behind their positive evaluation and explore how our system could be integrated into real-world car design workflows.

Easy to learn and use. The designers quickly learned and used 3D sketching and AI-based 2D rendering, thanks to their familiarity with perspective drawing and rendering, achieving satisfactory results after a one-hour tutorial. They strongly agreed the system was easy to learn and use, giving an average score of 4.7 (±0.3) on Q21-22 (Figure 14). The designers described the system as *"intuitive"* (P1, P4) and *"highly automated"* (P5), and noted that 3D sketching felt *"almost like sketching on paper"* (P1, P2, P3), and that *"any designer familiar with 3D tools would quickly get used to it"* (P3).

Highly compatible with existing tools. The designers appreciated the high level of compatibility between our system and professional tools widely used in the field of car design. They noted *"the system integrates seamlessly with industry-standard tools like Photoshop and Vizcom, making it easy to adopt as a complementary tool for current workflows"* (P6). For example, because designers controlled color with text prompts, tones did not always match exactly. Nevertheless, they could immediately open images in Photoshop, make quick adjustments using tools such as *Replace Color* or *Hue/Saturation* within seconds, and return to the system to continue their work without interruption. Moreover, the designers emphasized that the system's 3D sketch and 2D generated image outputs can be *"smoothly imported"* (P1, P5) into common 3D modeling software such as *"Alias or Blender"* (P5) as intermediate deliverables. Thanks to this compatibility, the designers anticipated that integrating the system into standard industry workflows would cause little friction.

Versatile. The designers highlighted several ways our system could be beneficial. First, in the ideation phase, they noted that it would be useful not only for *"quickly exploring"* (P1, P3, P4) a wide range of ideas but also for *"refining a single idea"* (P3, P4, P5, P6) into a cohesive three-view rendering set. Additionally, designers pointed out that the system naturally produces not only 2D renderings but also corresponding 3D curves and meshes, which could serve as valuable references in subsequent processes. They anticipated that these features would accelerate communication during both early-stage team reviews (P1, P2, P3, P4, P5) and later engineering phases (P3, P4, P5, P6). Some specifically mentioned that the system could be particularly useful for *"facelift projects"* (P2, P4).

7 LIMITATIONS & FUTURE WORK

In this section, we discuss the limitations identified in our user study, along with opportunities for future research.

Embracing 3D meshes as first-class citizens. In this study, 3D meshes primarily served a supporting role to 3D sketching. However, elevating 3D meshes to a primary representation, alongside 3D sketches, 2D drawings, and 2D renderings, could enable more flexible and diverse design workflows. To achieve this, generative AI models capable of creating 3D meshes from 2D sketches [12] could be integrated, along with interaction techniques that enable intuitive mesh editing through pen gestures [22, 42, 43].

Propagating changes across representations. In this study, the designers had to take additional steps to align one representation after modifying the other. Automating this process would allow designers to focus more fully on creative decision-making. For example, it is common practice for car designers to annotate on 2D renderings using distinctly colored pens during design reviews. Generative AI could reflect such feedback in the 2D image and automatically update the 3D sketch and 3D mesh accordingly.

Going VR. By placing 3D sketches and meshes at real scale within VR environments and projection-mapping 2D renderings onto meshes, designers could conduct immersive design reviews more quickly and affordably compared to creating high-fidelity 3D CAD models. This capability could significantly accelerate early-stage concept development. In addition, enabling designers to sketch 3D curves on a physical graphics tablet in VR [14, 32, 33] or directly manipulate 3D curves [34] or surfaces [35, 57] with hand gestures could further innovate car exterior design workflows, which still rely heavily on physical clay modeling for real-scale design iteration.

Expanding to other domains. While this study focused on validating the system's effectiveness within car exterior design, the need for view-consistent 2D renderings from multiple viewpoints exists in many design fields. We believe the proposed workflow is broadly applicable to various modes of transportation, both real and fictional (Figure 20), and plan to further enhance the system's generalizability and flexibility through collaborations with designers working in these diverse fields.

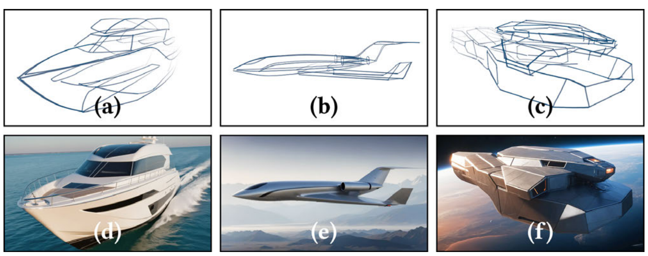

Figure 20: One of the authors used the proposed system to 3D sketch a variety of mobility types beyond cars, including (a) yachts, (b) airplanes, and (c) spacecraft. Based on these 3D sketches, the author was able to generate renderings in (d) front three-quarter, (e) side, and (f) rear three-quarter views.

8 CONCLUSION

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.

We conducted a user study with six professional car designers averaging 7.5 years of work experience to evaluate the practicality and effectiveness of the system. Over a total of 16 hours, the designers successfully created exterior design concepts for all 12 major car body types using the system. Results showed that the designers produced 12 refined 3D sketches and 36 high-quality 2D renderings approximately four to five times faster than conventional methods, all while maintaining a high level of view-consistency.

Designers used 3D sketching to express core forms and retain creative control throughout the process rather than delegating it to AI, progressively developing their designs through three distinct patterns: single-rounder, multi-rounder, and continuous switcher. Thanks to the tightly integrated system, designers navigated smoothly between 2D and 3D tasks with a high level of immersion, and maintained a clear overview of their design process. They responded positively to the prospect of the system being adopted within professional car design workflows.

We believe that in the era of generative AI, 3D sketching can play a central role in maximizing designers' creativity and productivity. In particular, the ability to generate an unlimited number of view-consistent 2D images from a single 3D sketch is expected to become increasingly valuable during the early stages of car design. We look forward to bringing further innovation to car exterior design and other industrial design processes through the addition of 3D generative AI and immersive VR workspaces.

Acknowledgments

This research was supported by the DRB-KAIST SketchTheFuture Research Center, the KAIST Convergence Research Institute Operation Program, and the InnoCORE Program of the Ministry of Science and ICT (N10250155). We thank Jaeseob Yeon, Jamin Koo, Jung Hwan Kwak, Sungguk Park, Kyung Han Chung, and Minchul Kim for participating in the car design workshop. We also thank Siripon Sutthiwanna and Hyelim Hwang for their technical support.

References

[1] Fatemeh Abbasinejad, Pushkar Joshi, and Nina Amenta. 2012. Surface patches from unorganized space curves. In *Proc. SoCG '12*. 417–418.
[2] Fatemeh Abbasinejad, Pushkar Joshi, Cindy Grimm, Nina Amenta, and Lance Simons. 2013. Surface patches for 3D sketching. In *Proc. SBIM '13*. 53–60.
[3] Maneesh Agrawala, Andrew C. Beers, and Marc Levoy. 1995. 3D painting on scanned surfaces. In *Proc. I3D '95*. 145–151.
[4] Erhan Batuhan Arisoy, Gunay Orbay, and Levent Burak Kara. 2012. Free form surface skinning of 3d curve clouds for conceptual shape design. *J. Comput. Inf. Sci. Eng.* 12, 3 (2012), 13 pages.
[5] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2008. ILoveSketch: as-natural-as-possible sketching system for creating 3D curve models. In *Proc. UIST '08*. 151–160.
[6] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2009. EverybodyLovesSketch: 3D sketching for a broader audience. In *Proc. UIST '09*. 59–68.
[7] Seok-Hyung Bae, Ryugo Kijima, and Won-Sup Kim. 2003. Digital styling for designers: 3D plane-symmetric freeform curve creation using sketch interface. In *Proc. ICCSA '03*. 701–710.
[8] Mikhail Bessmeltsev, Caoyu Wang, Alla Sheffer, and Karan Singh. 2012. Design-driven quadrangulation of closed 3D curves. *Trans. Graph.* 31, 6, Article 178 (2012), 11 pages.
[9] Carole Bouchard, Améziane Aoussat, and Robert Duchamp. 2006. Role of sketching in conceptual design of car styling. *J. Des. Res.* 5, 1 (2006), 116–148.
[10] Bill Buxton. 2007. *Sketching User Experiences.* Morgan Kaufmann.
[11] Liuqing Chen, Qianzhi Jing, Yixin Tsang, Qianyi Wang, Ruocong Liu, Duowei Xia, Yunzhan Zhou, and Lingyun Sun. 2024. AutoSpark: supporting automobile appearance design ideation with Kansei engineering and generative AI. In *Proc. UIST '24*. 1–19.
[12] 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
[13] European Commission. 2025. EU classification of vehicle types. https://alternative-fuels-observatory.ec.europa.eu/general-information/vehicle-types. Accessed: 2025-04-08.
[14] Tobias Drey, Jan Gugenheimer, Julian Karlbauer, Maximilian Milo, and Enrico Rukzio. 2020. VRSketchIn: exploring the design space of pen and tablet interaction for 3D sketching in virtual reality. In *Proc. CHI '20*. 1–14.
[15] Xusheng Du, Ruihan Gui, Zhengyang Wang, Ye Zhang, and Haoran Xie. 2025. Multi-view depth consistent image generation using generative AI models: application on architectural design of university buildings. arXiv:2503.03068

[16] Fabio Filippini and Gabriele Ferraresi. 2021. *Curve: 15 Lezioni sul Car Design.* Rizzoli Lizard.

[17] Tovi Grossman, Ravin Balakrishnan, Gordon Kurtenbach, George Fitzmaurice, Azam Khan, and Bill Buxton. 2002. Creating principal 3D curves with digital tape drawing. In *Proc. CHI '02.* 121–128.

[18] Pat Hanrahan and Paul Haeberli. 1990. Direct WYSIWYG painting and texturing on 3D shapes. *SIGGRAPH Comput. Graph.* 24, 4 (1990), 215–223.

[19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Proc. NeurIPS '20.* 6840–6851.

[20] Fei Hou, Chiyu Wang, Wencheng Wang, Hong Qin, Chen Qian, and Ying He. 2022. Iterative poisson surface reconstruction (iPSR) for unoriented points. *Trans. Graph.* 41, 4, Article 128 (2022), 13 pages.

[21] Zhiyang Huang, Nathan Carr, and Tao Ju. 2019. Variational implicit point set surfaces. *Trans. Graph.* 38, 4, Article 124 (2019), 13 pages.

[22] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 1999. Teddy: a sketching interface for 3D freeform design. In *Proc. SIGGRAPH '99.* 409–416.

[23] Kiia Kallio. 2005. 3D6B editor: projective 3D sketching with line-based rendering. In *Proc. SBIM '05.* 73–79.

[24] Levent Burak Kara and Kenji Shimada. 2007. Sketch-based 3D-shape creation for industrial styling design. *Comput. Graph. Appl.* 27, 1 (2007), 60–71.

[25] Olga Karpenko, John F. Hughes, and Ramesh Raskar. 2004. Epipolar methods for multi-view sketching. In *Proc. SBM '04.* 167–173.

[26] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. 2006. Poisson surface reconstruction. In *Proc. SGP '06.* 61–70.

[27] Michael Kazhdan and Hugues Hoppe. 2013. Screened poisson surface reconstruction. *Trans. Graph.* 32, 3, Article 29 (2013), 13 pages.

[28] 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.

[29] Yongkwan Kim and Seok-Hyung Bae. 2016. SketchingWithHands: 3D sketching handheld products with first-person hand posture. In *Proc UIST '16.* 797–808.

[30] Subhadeep Koley, Ayan Kumar Bhunia, Deeptanshu Sekhri, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. 2024. It's all about your sketch: democratising sketch control in diffusion models. In *Proc. CVPR '24.* 7204–7214.

[31] Joon Hyub Lee, Hanbit Kim, and Seok-Hyung Bae. 2022. Rapid design of articulated objects. *Trans. Graph.* 41, 4, Article 89 (2022), 8 pages.

[32] Joon Hyub Lee, Hyunsik Oh, Junwoo Yoon, Seung-Jun Lee, Taegyu Jin, Jemin Hwangbo, and Seok-Hyung Bae. 2024. RobotSketch: a real-time live showcase of superfast design of legged robots. In *SIGGRAPH Asia Real-Time Live! '24.* Article 9, 1 pages.

[33] 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 *SIGGRAPH Emerging Technologies '24.* Article 17, 2 pages.

[34] Sang-Hyun Lee, Taegyu Jin, Joon Hyub Lee, and Seok-Hyung Bae. 2022. WireSketch: bimanual interactions for 3D curve networks in VR. In *UIST '22 Adjunct.* Article 37, 3 pages.

[35] Sang-Hyun Lee, Joon Hyub Lee, and Seok-Hyung Bae. 2024. Bimanual interactions for surfacing curve networks in VR. In *Proc. CHI EA '24.* Article 60, 7 pages.

[36] 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 *Proc. CHI EA '25.* Article 352, 7 pages.

[37] Tony Lewin and Ryan Borroff. 2010. *How to Design Cars Like a Pro.* Motorbooks.

[38] David Chuan-En Lin, Hyeonsu B Kang, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K Hong. 2024. Inkspire: sketching product designs with AI. In *Proc. UIST '24.* Article 86, 6 pages.

[39] Vivian Liu, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2023. 3DALL-E: integrating text-to-image AI in 3D design workflows. In *Proc. DIS '23.* 1955–1977.

[40] Stuart Macey and Geoff Wardle. 2014. *H-Point: The Fundamentals of Car Design & Packaging.* Design Studio Press.

[41] 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

[42] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. 2007. FiberMesh: designing freeform surfaces with 3D curves. *Trans. Graph.* 26, 3 (2007), 41–50.

[43] Andrew Nealen, Olga Sorkine, Marc Alexa, and Daniel Cohen-Or. 2005. A sketch-based interface for detail-preserving mesh editing. *ACM Trans. Graph.* 24, 3 (2005), 1142–1147.

[44] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. 2022. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. arXiv:2112.10741

[45] Jeongseok Oh, Seungju Kim, and Seungjun Kim. 2024. LumiMood: a creativity support tool for designing the mood of a 3D scene. In *Proc. CHI '24.* Article 174, 21 pages.

[46] Günay Orbay and Levent Burak Kara. 2012. Sketch-based surface design using malleable curve networks. *Comput. Graph.* 36, 8 (2012), 916–929.

[47] Hao Pan, Yang Liu, Alla Sheffer, Nicholas Vining, Chang-Jian Li, and Wenping Wang. 2015. Flow aligned surfacing of curve networks. *Trans. Graph.* 34, 4, Article 127 (2015), 10 pages.

[48] Songyou Peng, Chiyu Jiang, Yiyi Liao, Michael Niemeyer, Marc Pollefeys, and Andreas Geiger. 2021. Shape as points: a differentiable poisson solver. In *Proc. NIPS '21.* Article 998, 13 pages.

[49] Moreno Attilio Piccolotto. 1998. *Sketchpad+: Architectural Modeling Through Perspective Sketching on a Pen-Based Display.* Master's thesis. Cornell University.

[50] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with CLIP latents. arXiv:2204.06125

[51] Gus Rick. 2025. Graphic artists in China push back on AI and its averaging effect. https://www.theverge.com/ai-artificial-intelligence/688645/graphic-artists-china-ai. Accessed: 2025-06-27.

[52] Scott Robertson and Thomas Bertling. 2013. *How to Draw: Drawing and Sketching Objects and Environments from Your Imagination.* Designstudio Press.

[53] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proc. CVPR '22.* 10684–10695.

[54] Bardia Sadri and Karan Singh. 2014. Flow-complex-based shape reconstruction from 3D curves. *Trans. Graph.* 33, 2, Article 20 (2014), 15 pages.

[55] 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. Photorealistic text-to-image diffusion models with deep language understanding. In *Proc. NeurIPS '22*, Vol. 35. 36479–36494.

[56] Yulin Shen, Yifei Shen, Jiawen Cheng, Chutian Jiang, Mingming Fan, and Zeyu Wang. 2024. Neural canvas: supporting scenic design prototyping by integrating 3D sketching and generative AI. In *Proc. CHI '24.* Article 1056, 18 pages.

[57] Gravity Sketch. 2017. Gravity Sketch.

[58] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising diffusion implicit models. arXiv:2010.02502

[59] Olga Sorkine, Daniel Cohen-Or, Yaron Lipman, Marc Alexa, Christian Rössl, and H.-P. Seidel. 2004. Laplacian surface editing. In *Proc. SGP '04.* 175–184.

[60] Tibor Stanko, Stefanie Hahmann, Georges-Pierre Bonneau, and Nathalie Saguin-Sprynski. 2016. Smooth interpolation of curve networks with surface normals. In *Proc. Eurographics '16.* 21–24.

[61] Udin. 2024. Optic.

[62] Vizcom. 2023. Vizcom.

[63] Qiang Wang, Di Kong, Fengyin Lin, and Yonggang Qi. 2023. DiffSketching: sketch control image synthesis with diffusion models. arXiv:2305.18812

[64] Min Xin, Ehud Sharlin, and Mario Costa Sousa. 2008. Napkin sketch: handheld mixed reality 3D sketching. In *Proc. VRST '08.* 223–226.

[65] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. 2024. Consistnet: enforcing 3D consistency for multi-view images diffusion. In *Proc. CVPR '24.* 7079–7088.

[66] Emilie Yu, Rahul Arora, J. Andreas Bærentzen, Karan Singh, and Adrien Bousseau. 2022. Piecewise-smooth surface fitting onto unstructured 3D sketches. *Trans. Graph.* 41, 4, Article 88 (2022), 16 pages.

[67] Emilie Yu, Rahul Arora, Tibor Stanko, J. Andreas Bærentzen, Karan Singh, and Adrien Bousseau. 2021. Cassie: curve and surface sketching in immersive environments. In *Proc. CHI '21.* 1–14.

[68] Hongbo Zhang, Pei Chen, Xuelong Xie, Zhaoqu Jiang, Yifei Wu, Zejian Li, Xiaoyu Chen, and Lingyun Sun. 2025. FusionProtor: a mixed-prototype tool for component-level physical-to-virtual 3D transition and simulation. In *Proc. CHI '25.* Article 685, 19 pages.

[69] Hongbo Zhang, Pei Chen, Xuelong Xie, Chaoyi Lin, Zhuoshu Li Lianyan Liu, Weitao You, and Lingyun Sun. 2024. ProtoDreamer: a mixed-prototype tool combining physical model and generative AI to support conceptual design. In *Proc. UIST '24.* Article 97, 18 pages.

[70] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. In *Proc. ICCV '23.* 3836–3847.

[71] Yixin Zhuang, Ming Zou, Nathan Carr, and Tao Ju. 2013. A general and efficient method for finding cycles in 3D curve networks. *Trans. Graph.* 32, 6, Article 180 (2013), 10 pages.