Aesthetic3D: Incorporating Shape Aesthetic Measures into 3D Modeling Interfaces

DENG YU and JIANING GUO, Shandong University, China and Qingdao Key Laboratory of Trustworthy Artificial Intelligence, China
HUI YE, Hong Kong Baptist University, China
PENGJIE REN*, Shandong University, China
HONGBO FU*, Hong Kong University of Science and Technology, China
MANFRED LAU*, Lakehead University, Canada

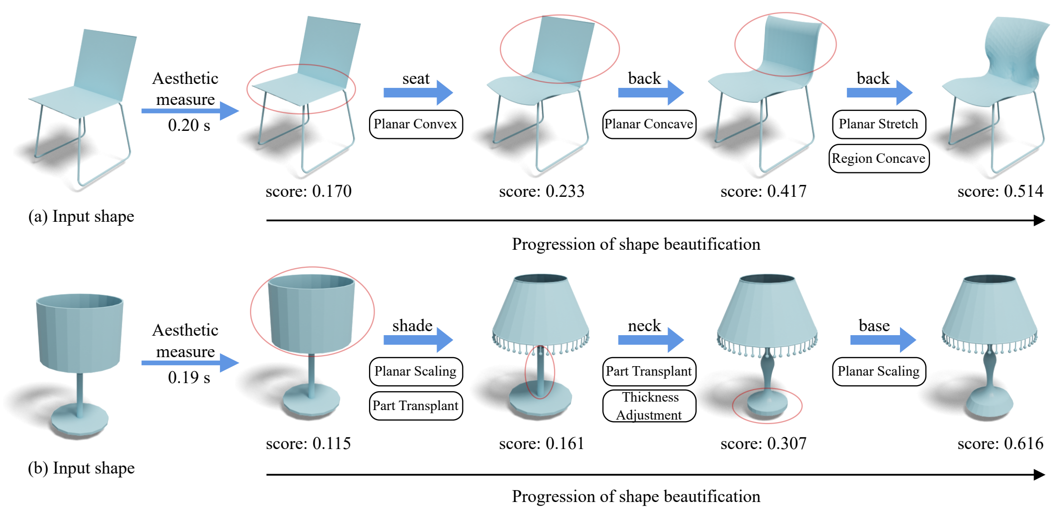

Fig. 1. Progression of shape beautification with our *Aesthetic3D* work. Given an input shape (a) or (b), *Aesthetic3D* can compute aesthetic scores for them at real-time speed, and provide a suite of aesthetics-driven editing tools (see example tools under each arrow), letting users effectively adjust 3D shapes based on the computed scores until the desired aesthetic effects are achieved.

3D Aesthetics is significant in digital design, shaping how users experience real-time 3D content in games, VR, and product design. However, creating aesthetically pleasing shapes remains challenging due to diverse subjective standards and the lack of tools that support aesthetics-driven editing. Users often rely on intuition without explicit guidance on visual appeal, making aesthetics refinement slow, inconsistent, and cognitively demanding, particularly in fast-paced, iterative workflows. To address this challenge, we conducted in-depth

*Corresponding authors

Authors' Contact Information: Deng Yu, deng.yu@sdu.edu.cn; Jianing Guo, Shandong University, China and Qingdao Key Laboratory of Trustworthy Artificial Intelligence, China; Hui Ye, huiyehy@hkbu.edu.hk, Hong Kong Baptist University, China; Pengjie Ren, Shandong University, China; Hongbo Fu, fuplus@gmail.com, Hong Kong University of Science and Technology, China; Manfred Lau, manfred.lau@lakeheadu.ca, Lakehead University, Canada.

ACM 2577-6193/2026/5-ART9
https://doi.org/10.1145/3804501

interviews with design experts to identify challenges in aesthetics-oriented modeling workflows. Based on the findings, we developed *Aesthetic3D*, a 3D modeling interface that provides real-time aesthetics scores learned from human perceptual data. Furthermore, *Aesthetic3D* seamlessly integrates the learned aesthetics measures into intuitive editing operations, enabling aesthetics-driven exploration and refinement of shape geometry. We evaluated *Aesthetic3D* through an ablation study, an open-ended study, and three generalization evaluations. Comprehensive experiments show that with *Aesthetic3D*, users can easily and effectively enhance the aesthetics appeal of 3D shapes.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; **Perception**; • **Human-centered computing** → **Graphical user interfaces**.

Additional Key Words and Phrases: 3D Shape Beautification, Aesthetics Guidance, Neural Aesthetics Measure, 3D Modeling Interface

ACM Reference Format:
Deng Yu, Jianing Guo, Hui Ye, Pengjie Ren, Hongbo Fu, and Manfred Lau. 2026. *Aesthetic3D*: Incorporating Shape Aesthetic Measures into 3D Modeling Interfaces. *Proc. ACM Comput. Graph. Interact. Tech.* 9, 1, Article 9 (May 2026), 20 pages. https://doi.org/10.1145/3804501

1 Introduction

Aesthetics plays a pivotal role in how people perceive, evaluate, and engage with digital 3D content across industrial design [6, 7, 13], animation [33, 59], and digital art [47, 56, 57]. In these fields, users strive to create models that are not only functionally effective but also visually pleasing. Unlike tangible characteristics such as geometry, topology, or texture, which are quantifiable and well-supported by established 3D modeling tools such as Autodesk Maya [1], Cinema 4D [34], ZBrush [35], etc., characterizing shape aesthetics during the 3D modeling process remains largely subjective and under-supported. To the best of our knowledge, no existing tools provide the explicit guidance or feedback on the visual appeal of the objects that are being modeled during the interactive 3D modeling process. Therefore, designers have to rely heavily on intuition, experience, and iterative trial-and-error when refining 3D shapes to make them aesthetic. This hinders fast-paced creative workflows that demand rapid iteration and decision-making.

Recent advances in computational aesthetic measurement [12, 13, 17] offer promising opportunities to bridge this gap. Currently, most algorithms are designed for the aesthetic beautification of 2D images [16, 22, 53, 60], such as layout optimization [53] and photo retouching [60]. For 3D shapes, prior works [4, 38, 42–44] build relationships between aesthetics properties suggested by art and philosophy with manually-defined geometric features (e.g., curvature, symmetry, proportion) or mathematical criteria (e.g., bending energy, minimum variation surface). However, manually-defined features are difficult to generalize across different shape styles and user preferences. More recent techniques [12, 13, 17] leverage data-driven approaches to learn the user preferences of shape aesthetics directly from shape aesthetics datasets. However, prior approaches lack methods for guiding users about how their design choices affect the perceived aesthetics of the 3D shape that is being modeled.

In order to investigate the key challenges and underlying needs in the design workflows of aesthetics-oriented 3D modeling scenarios, we conduct a formative study by interviewing expert designers experienced in interior design and furniture design. Based on the interview findings, we summarize their motivations, methods, difficulties in designing and editing 3D shapes, and their expectations for an ideal shape beautification tool. We then formulate four design considerations for such an ideal system (Figure 1): *(i) the use of aesthetics measures during the 3D modeling process, (ii) aesthetics-oriented 3D editing in a coarse-to-fine manner, (iii) iterative re-scoring and re-designing, (iv) interactive scoring-and-editing at real-time speed.*

Based on the design considerations, we present *Aesthetic3D*, an interactive 3D modeling interface with aesthetics scoring guidance, enabling users to iteratively refine shapes by providing continuous, real-time feedback on the aesthetic quality of their designed shapes. *Aesthetic3D* consists of two key components: (1) a real-time perceptual aesthetics scoring module and (2) a suite of aesthetics-guided editing tools. The former is learned from the human-annotated aesthetics ratings and computes a real-time aesthetics score for a given shape. The latter are designed in a coarse-to-fine scale (based on the editing effects on 3D shapes) that mimics common designers' workflows. This allows *Aesthetic3D* to help users obtain immediate visual and quantitative feedback on the shape aesthetics during the modeling process, making it easier for users to create more complex shapes.

To evaluate *Aesthetic3D*, we conduct an ablation study, an open-ended study, and three generalization evaluations among both experts in art and design and novice users. From the analysis of their modeling processes, the comparisons of their diverse beautified results, and discussions with them, we find that *Aesthetic3D* is an easy-to-use and efficient modeling tool that can help users create aesthetic 3D shapes with highly useful guidance.

In summary, the main contributions of our work can be summarized as follows:

- We identify the challenges and needs of users in aesthetics-oriented 3D modeling from in-depth expert interviews, and use these findings to guide the design of our system.
- We introduce an intuitive aesthetics-guided 3D shape editing approach that allows users to beautify 3D shapes with their aesthetics scores computed and displayed interactively.
- We incorporate the 3D shape aesthetics measures into 3D modeling interfaces, and present a suite of 3D shape editing tools in a coarse-to-fine scale to facilitate the modeling of aesthetic 3D shapes.
- We validate the usability and effectiveness of the interface through user studies, and show diverse and aesthetic 3D models created with the system.

2 Related Work

2.1 Computational Aesthetic Measures

While aesthetics is a highly abstract and theoretical concept in philosophy and art, researchers in computer graphics and vision endeavour to investigate computational metrics for the explicit evaluation of aesthetics. In the literature, existing works in computational aesthetic measures can be broadly categorized into two branches: image-based [5, 27, 45] and shape-based [17, 38].

For 2D images, early studies [15, 26, 39] relied on handcrafted features describing visual and compositional attributes to explicitly model well-established photographic rules. More recent works [27, 36, 54] focused on end-to-end learning of aesthetics, facilitated by advances of neural networks and the growth of large-scale annotated image datasets. However, supervised by the aforementioned aesthetic metrics or aesthetics-revealing features, 2D image beautification remains confined to color enhancement [2, 16, 53], composition adjustment [20, 22, 28], and aesthetic image generation [7], thus are inadequate for 3D shapes involving complex geometry editing and structure manipulation.

In the 3D domain, early works [42, 43] explored the relationships between aesthetic characteristics and linguistic or parametric design variables in 3D designs via empirical investigations. Following this, mathematical models [4, 38] for 3D aesthetic design were introduced, extending 2D aesthetic criteria to 3D geometric features [4] or formulating aesthetic curves and surfaces with well-defined splines [38]. More recently, Dev and Lau [17] proposed a cross-category end-to-end 3D geometry aesthetics assessment trained on their curated shape aesthetic dataset. Using the same dataset, a subsequent work [12] proposed a patch-based learning framework to further predict both global aesthetic scores and local aesthetic maps for 3D shapes. Other efforts explored aesthetic enhancement for specific contexts, such as 3D layouts [52], 3D sketch strokes [32], etc., by predefined geometric and spatial constraints related to symmetry, angles, proportions, and alignment. However, such task-specific aesthetic measures do not generalize well to broader 3D modeling scenarios.

The most relevant work to ours is [13], where Chen and Lau proposed a reference-guided approach to deform an unattractive input shape into its more beautiful reference shape retrieved from the dataset. However, the lengthy optimization time of tens of minutes and the reference-guided approach are limited by dataset size. Furthermore, it struggles when the input and reference differ significantly, hindering the practical usage of their method. *Aesthetic3D* fills this gap by (1) integrating 3D shape aesthetic measures into interactive design that supports real-time response, and (2) replacing the reference-guided approach with a more direct aesthetic score–guided way.

2.2 Aesthetics-oriented 3D Modeling Interfaces

Aesthetic considerations, such as symmetry, curvature, balance, and proportion, span the entire lifecycle of design, but most commercial modeling systems, such as Autodesk Maya [1], Cinema 4D [34], etc., offer little to no support for evaluating or enhancing a shape's visual appeal.

Existing works [11, 25, 51, 55] for aesthetics typically tackle this challenge indirectly, falling into two general categories: either by employing style transfer techniques [24, 55] or by retrieving parts and shapes from aesthetically pleasing examples [11, 25]. Following this, various data-driven modeling frameworks [13, 40, 51] have incorporated examples of high-quality shapes to guide synthesis or recombination. However, these methods are generally offline and lack interactive feedback, thus cannot inform users about aesthetic quality in real time. Our work takes a further step in advancing this line of research by introducing *Aesthetic3D*, the first interactive 3D modeling system (to our knowledge) that couples learned aesthetic measures with real-time visual feedback.

2.3 3D Shape Editing

The advancement of 3D shape editing techniques [14, 30, 37, 58] has significantly expanded the computational expressiveness and interaction paradigms of aesthetic-enhancement systems. Depending on the employed shape representations, existing research can be broadly divided into two major classes: *implicit*-representation editing [37, 41, 58], which offers topological flexibility and smooth continuity, and *explicit* geometric editing [14, 30], which provides direct control over surface structure and fine-grained geometry. Note that a comprehensive review of shape generation and shape editing methods is beyond the scope of this work.

Implicit representations such as DeepSDF [41], Occupancy Networks [37], and Neural Radiance Fields [58] underpin differentiable editing frameworks. For example, leveraging continuity and topological flexibility of these fields, methods like NeRF-Editing [58] and Instruct-NeRF2NeRF [21] enable controlled modification of appearance and geometry through conditional pose-level or language-guided, semantics-level editing. However, lacking fine-grained geometric controllability and real-time response, this line of methods is less suited for our interactive aesthetic enhancement.

Explicit geometric editing methods [14, 30] operate directly on meshes, point clouds, and polygonal surfaces, enabling fine-grained shape manipulation across multiple scales. For example, GenPara [14] supports high-fidelity structural editing via multi-granularity selection of local point sets or fragments, while MeshDiffusion [30] and its variants exploit the tetrahedral parametric way. Compared with implicit approaches, explicit approaches offer stronger local controllability and better interpretability, aligning well with interactive design needs. Our system also builds on explicit representations and corresponding editing operations to support direct, real-time, and interactive shape editing with aesthetic guidance.

3 Formative Study

To better understand aesthetics-oriented design workflows in practical 3D modeling scenarios and to identify the challenges involved, we conducted an in-depth semi-structured interview study.

3.1 Interviewees

We recruited five interviewees through our personal and research networks. The interviewees consisted of 5 design experts (E1-E5, 5 males): 1 interior designer (E1), 3 furniture designers (E2, E3, E4), and 1 3D modeler (E5). They had 5-20 years (mean: 11 years, SD: 4.90 years) experience in furniture design, scene modeling or interior decoration. They worked closely with design studios to provide user-customized and aesthetically pleasing products, tailored proposals, and services (e.g., aesthetics trend forecasting, brand aesthetic positioning, and brand aesthetic heritage).

3.2 Interview Protocol

We conducted interviews remotely via Zoom video/audio conferencing tool. After collecting their background information, we asked them about their 3D modeling tools for design and workflows, common challenges and problems during 3D shape aesthetics editing, and expected features of a novel tool. They were allowed to freely talk about their experiences and opinions during a 30-40 minute interview. Please refer to the supplementary materials for the detailed interview questions.

3.3 Data Analysis

We used an open coding method [9, 10] to collect and analyze the interview data. Two authors conducted the coding process. We discussed all the codes among the co-authors to identify emerging themes, which were then grouped into two themes (Sections 3.4.1 and 3.4.2). Within the second theme, we further sub-categorized the codes into several sub-themes. We sorted these sub-themes by mentioned frequency and distilled five main issues (I1-I5), summarized as follows.

3.4 Findings

3.4.1 Conventional design workflow and methods. All participants typically design 3D shapes in a coarse-to-fine manner. They typically first blocked out a rough convex hull for a target shape using simple primitives (e.g., cubes or spheres) and then refine local details towards the desired form, with their years of design experience. To produce qualified 3D shapes, they relied on professional software (e.g., Autodesk Maya, 3ds Max, Cinema 4D, etc.) to create 3D shapes and effects in the modeling stage (E1, E2, E3, and E5) for design workflows and scenarios (E2, E3, and E4). Becoming a competent industry designer, typically required new starters to review and analyze large number of online product samples diverse aesthetic styles (E3 and E4). They aimed to communicate and evaluate designed ideas by rapidly modeling 3D products regarding the aesthetic appeal (E1-E5).

3.4.2 Common Difficulties and Expectations. **I1: shape aesthetic measures.** All participants emphasized that measuring shape aesthetics is an essential dimension for product design. However, their individual aesthetic judgment can be difficult to persuade other cooperators, leaders, or customers, thus hampering effective communication and evaluation of design ideas. E3 mentioned that "*disagreements on aesthetics arise frequently (also reported by E1, E2, E4 and E5) and are often resolved only by repeatedly revisiting higher design concepts or original goals to reach consensus*". This further incurs higher time costs and additional workload.

All participants reported a broad consensus on aesthetic aspects of shape, such as the golden ratio and overall harmony. E4 described a distinction between **hard** and **soft** aesthetic standards: "*hard standards are rule-based (e.g., curvature magnitudes) and precisely computable, whereas soft standards are perception-based and can be coarsely estimated by ranges or thresholds.*" According to

E4, hard standards: 1) are often partially defined, 2) may conflict with others (e.g., symmetry *vs.* vitality sense), 3) may be optimal in one design context but undesired in another. While historically soft standards such as the golden ratio (mentioned by E1–E3, E5) can become widely accepted and eventually solidify into hard standards through broad adoption. E1 and E4 further pointed out that "*establishing new soft aesthetic standards is expensive, since it typically requires extensive expert involvement for rating and labeling*". Despite this cost, all participants preferred soft standards because they are more widely defined and better aligned with common human perception.

I2: coarse-to-fine workflow. All participants reported using a coarse-to-fine design manner in their modeling workflows but found it inconvenient to perform coarse and fine-grained edits when it required frequently switching between different software. E2 mentioned "*3ds Max works well on global shaping but is unfriendly for refining curved local regions. Therefore, I need to switch to Maya for local design and then back again, which is cumbersome and tedious.*" It resonates with E5's comments "*it is important to have a platform integrated coarse-to-fine editing tools for aesthetics-oriented 3D modeling. It would allow faster prototyping and help me focus more on ideations and creations*".

I3: freeform aesthetics-driven editing tools. All participants mentioned that mainstream 3D software rarely provides aesthetics-oriented editing functions or tools. This confirms the challenges identified in the previous works, i.e., "*while the aesthetics of images has been extensively researched to date, work in 3D shape aesthetics is very limited*" in [17]. E1, E3, E4, and E5 reported that "*it's challenging for traditional 3D software to perform creative 3D modeling or shape beautification because everything is quantified by parameters, unfriendly for freeform modification, fast embodiment, or smooth beautification*". All participants advocated for more non-parametric and intuitive editing capabilities to better support aesthetics-relevant 3D modeling in a novel tool.

I4: re-scoring and re-design. Participants E1 and E5 mentioned that aesthetics-oriented 3D modeling was supposed to be both reversible and traceable. E1 further emphasized that "*beautification should not be a destruction process, we need to preserve the prior designed versions in case designers go back to the earlier state*". In this case, designers tend to restore the original shape and withdraw the operations. Without considering this, the whole process will become more time-consuming.

I5: real-time feedback. For a system supporting aesthetics-guidance 3D modeling, the participants believed that aesthetic guidance (scores) should respond to the user's edits as quickly as possible. E1 and E5 emphasized that "*real-time feedback is also a very important aspect for designers in the 3D modeling process*". Although the real-time issue was not emphasized by all participants, we treat it as a necessary design choice and a default goal in our system.

According to these findings, we have four important design considerations for our *Aesthetic3D*:
DC1: incorporate common 3D shape aesthetic measures into modeling process.
DC2: support coarse-to-fine level aesthetics-oriented 3D editing.
DC3: support iterative re-score and re-design of 3D shapes easily.
DC4: support interactive scoring-and-editing at real-time speed.

4 System Design

Considering flexibility, compatibility, and ease of development for a freelance 3D system, we deploy *Aesthetic3D* on a Microsoft Surface with a stylus pen (Figure 3), built on the Blender framework. In the following, we introduce modeling walk-through (Section 4.1) and key components (Section 4.2).

4.1 Modeling Walk-through

Under the interactive mechanism, the usage of *Aesthetic3D* for 3D shape beautification can be divided into four primary steps: (1) object import and aesthetics scoring, (2) shape editing with tools, (3) aesthetics re-scoring and re-editing, and (4) iterative refinement (Figure 2 (a-d)).

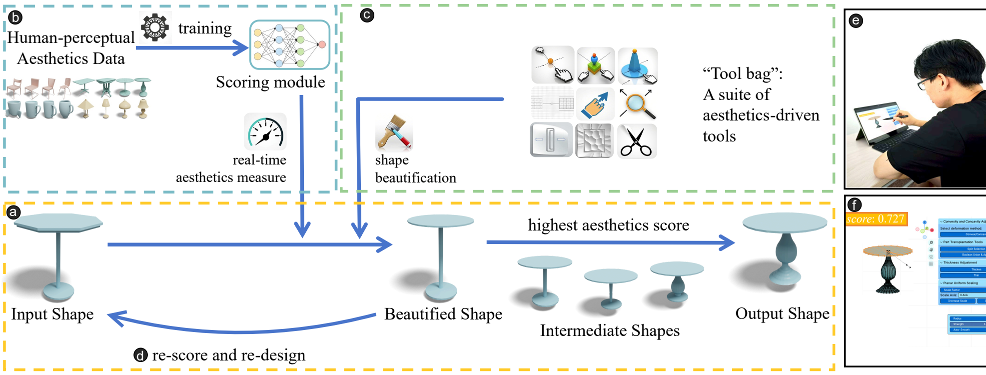

Fig. 2. Modeling walk-through of *Aesthetic3D* (a-d) and current design (e-f). During modeling (a), aesthetic scoring module (b) and aesthetics-driven editing tools (c), make aesthetic enhancement perceptible, measurable, editable, and iterative (d). The user (e) can easily adjust the aesthetic appeal of a 3D object by scaling the table surface with a stylus pen on our UI. (f) shows the screenshot of a designed shape in editing mode.

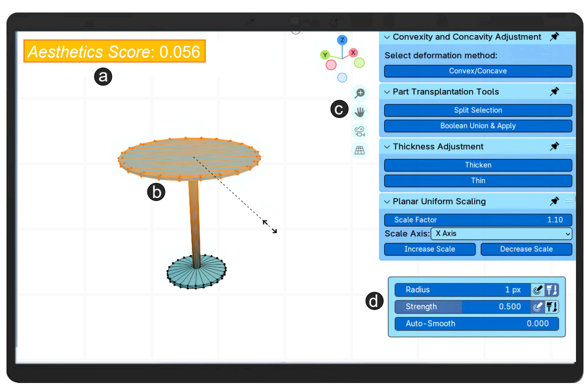

Fig. 3. Our Graphical Interface, presents four key elements: (a) computed aesthetic score of the current 3D shape, (b) being edited regions (yellow color), (c) a list of aesthetics-driven editing tools, and (d) a panel of three sliders for fine-grained control of editing effects, e.g., adjust the effective region size via "*Radius*" slider.

Step 1: object import and aesthetics scoring. When importing a 3D shape, it is first placed at the center of the interface. Building upon the Blender framework, our interface also supports two modes: object and edit modes. In object mode, user can view and manipulate the 3D object as a whole, while edit mode parses the 3D object into a mesh of triangular faces, allowing directly adjusting vertices, edges, or faces. Users can switch between two modes by pressing "*Tab*". In object mode, the initial aesthetics score can be manually triggered by pressing ("*Ctrl*" + "*Alt*" + "*R*").

Step 2: shape editing with tools. After obtaining the initial aesthetic score of the imported 3D shape, users can freely manipulate global proportions, surface characteristics, or local geometry with our aesthetics-driven editing tools. For example, given the input shape (Figure 2 (a)), the user first replaces the original tabletop with a more round one by the "part transplant" tool, and then refines local geometry, such as the thickness and length of the table legs, by applying the "thickness adjustment" tool to leg regions and the "vertex dragging" tool to selected vertices, respectively.

Step 3: aesthetic re-scoring and re-editing. Throughout the 3D beautification process, our *Aesthetic3D* system displays real-time aesthetic scores on top of the interface. When users wish to

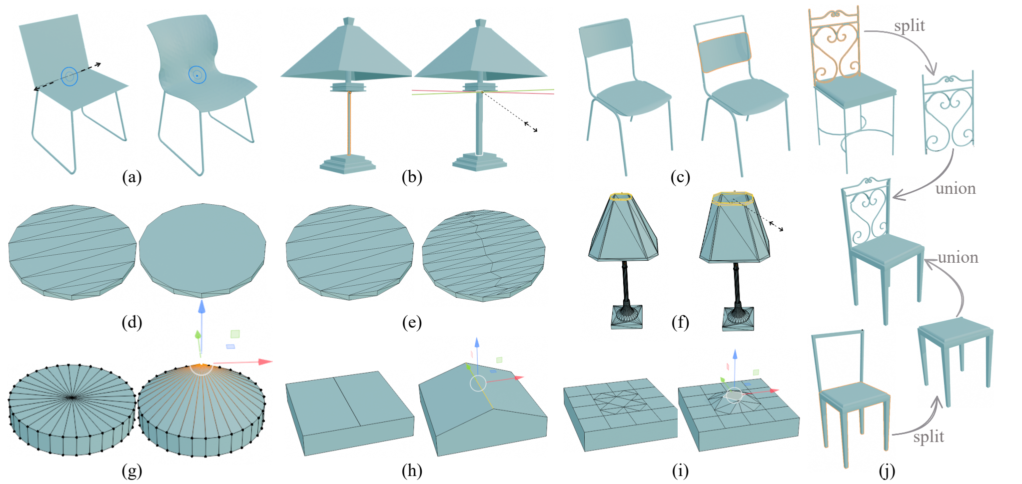

Fig. 4. Aesthetics-driven editing tools in *Aesthetic3D* support the following operations: (a) surface convexity and concavity adjustment, (b) region thickness adjustment, (c) region moving, (d-e) face subdivision or merge, (f) planar uniform scaling, (g) vertex dragging, (h) line dragging, (i) face dragging, and (j) part transplant.

explore alternative design directions or more complex aesthetic effects (see intermediate shapes in Figure 2), they can manually trigger scoring by pressing ("*Ctrl*" + "*Alt*" + "*R*") after completing an operation or automatically trigger real-time scoring by switching from edit mode to object mode. By comparing the scores before and after each edit, users can clearly measure the impact of their modifications and make informed design decisions.

Step 4: iterative refinement. Our system supports iterative refinements: if the aesthetic score increases after an edit, the user is encouraged to continue and optimize in that direction; otherwise, the user can undo the change and explore alternative strategies. Through repeated iterations, the edited 3D shape progressively converges toward the one that better fulfills the user's design intention with a higher aesthetic score.

4.2 Key Components

Based on the design considerations distilled from the interviews, we design two key components of *Aesthetic3D*: an aesthetic scoring module and a suite of aesthetics-driven editing tools.

Aesthetic scoring module. Our system enables users to explicitly measure the aesthetic score of a 3D object being designed (Figure 3 (a, b)), making the impact of editing operations on shape aesthetics observable throughout the modeling process (DC1). Once activated, this module runs in the background and updates the score when a new 3D shape is imported or an editing is completed.

To support real-time aesthetic scoring (DC4), our system: (1) renders a 3D shape into 12 multi-view images (uniformly sampled every 30 degrees in azimuth); (2) leverages EGL (Embedded Graphics Library) API to accelerate rendering time speed from around 1∼2 seconds to under 0.2 seconds; (3) feeds these multi-view images into a lightweight neural network that serves as the core neural aesthetic measure. Note that all the aforementioned functions run in the background. We elaborate our underlying neural aesthetic model in Section 5.

Aesthetics-driven editing tools. To support iterative, aesthetics-oriented scoring and designing of 3D shapes (DC3), our system further provides a suite of aesthetic editing tools (Figure 3 (c)) that allow users to repeatedly adjust their design until the desired aesthetic effect is achieved.

We summarize prior works [12, 13, 42, 43] and designers' suggestions (Section 3), and define 9 types of commonly-used aesthetics-driven editing tools (Figure 4), organized in a coarse-to-fine hierarchy (DC2). The Blender framework provides fundamental 3D editing functions such as vertex, line, and face operations, which partially reduces our coding burdens and allows us to focus more on the rest editing tools.

Detailed descriptions of these aesthetics-driven tools are provided in the supplementary materials. Our editing tools are structured hierarchically, ranging from global manipulations that alter the overall form to local refinements at the surface, line, and vertex levels. To enable more fine-grained user control, our system further provides interactive sliders (Figure 3 (d)). For example, users can easily adjust the affected region size and editing intensity by "*Radius*" and "*Strength*" sliders.

5 Neural Aesthetic Measure

In this section, we introduce the underlying algorithm of *Aesthetic3D* (scoring module), i.e., neural aesthetic model, from two aspects: data collection (Section 5.1) and learning strategy (Section 5.2).

5.1 Human Perceptual Aesthetic Data Collection

Inspired by prior works [12, 13, 17], we develop our aesthetic measure in a data-driven manner (Figure 2 (b)). However, collecting a large dataset of human-perceptual aesthetic annotations from design experts is challenging and costly. To mitigate the cost issue and enhance the generality of our learned aesthetic measures: (1) we first curate a subset of 3D shapes from ShapeNet [8], focusing on four representative categories: chairs, lamps, mugs, and tables; and (2) we invite both experts (3 artists) and ordinary people (13 graduate students in computer science) from our university to label shape data. We manually filtered out shapes with flaws (e.g., missing faces) in each category to ensure data quality. To further enrich the dataset diversity, we removed repeated or similar shapes from the filtered subset. Our final dataset comprises 277 chairs, 88 lamps, 74 mugs, and 40 tables.

Fig. 5. Representative human perceptual aesthetic data in our dataset. We show the collected shapes and the average rating for each shape.

To collect human-perceptual aesthetic annotations, we randomly divided shapes of a category into groups of five, with each group was independently rated by the participants from our institution. Participants could freely rotate and zoom each 3D model and rate their aesthetic agreement using a 5-point Likert scale, ranging from *strongly disagree* (-1) to *strongly agree* (1). Please examine the representative examples in our dataset in Figure 5.

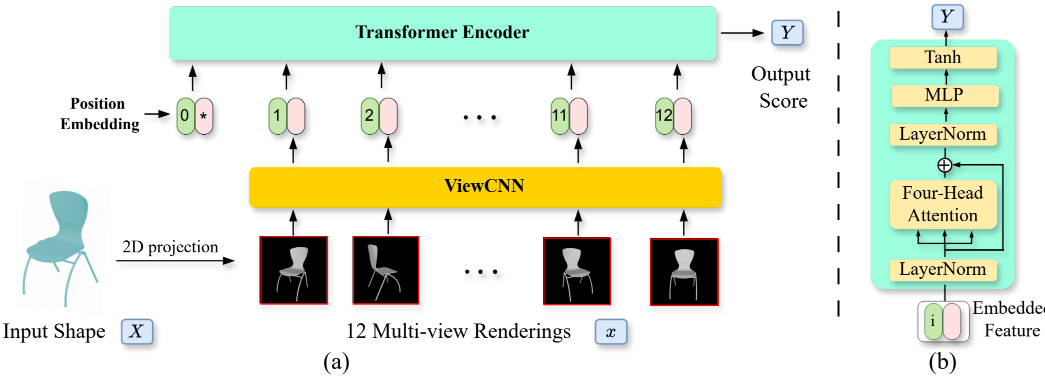

Fig. 6. The details of the neural aesthetic model for our *Aesthetic3D*.

5.2 Learning Aesthetic Measures

After obtaining human-perceptual aesthetic data (3D shapes X with corresponding aesthetic rating scores Y) and considering the real-time needs, we design a lightweight network $F_{W,b}$ (Figure 6) to learn the aesthetic mapping ($X \longrightarrow Y$) from training data.

As shown in Figure 6 (a), we represent a 3D shape X using 12 multi-view images x, uniformly sampled every 30 degrees in azimuth. Different from [17], our method employs a Vision Transformer (ViT) backbone (Figure 6 (a)) [18] to aggregate and encode features across multi-view images x.

$$Y = F_{W,b}(x) \tag{1}$$

More specifically, we first extract the view features (1×64) from each view image (256×256) using a simple CNN with four *Conv* and *Norm* layers (Figure 6 (a)), and then feed embedded features (13×128) (concatenate position tokens to view features) to the transformer (Figure 6 (b)), mapping the rendered multi-view images x of a 3D shape X to its corresponding aesthetic score Y.

To train the neural aesthetic model $F_{W,b}$, we adopt the same pairwise ranking strategy [17] to encourage our model to learn the ranking function over different input 3D shapes. More specifically, we randomly select multi-view images (x_A, x_B) of paired shapes (X_A, X_B) as training data and further minimize the loss function as below:

$$L(W, b) = \sum_{(x_A, x_B) \in I_{train}} \max\left(0, m - (Y_A - Y_B)\right) \tag{2}$$

where $Y_A = F_{W,b}(x_A)$, $Y_B = F_{W,b}(x_B)$, and $m = 1$.

6 Evaluation

We compare our learned neural aesthetic model with Dev and Lau's [17], since they have similar objectives to ours. We first split our collected dataset (Section 5.1) into training and test sets at a 7:1 ratio empirically, similar to [12], and then train two methods on the same training set, and finally report their performance (on test set) in Table 1.

Table 1. Accuracy of different methods on test set.

Methods	Chair	Table	Lamp	Mug	*Avg.* accuracy
Dev and Lau's [17]	0.409	0.355	0.401	0.384	0.387
Ours	**0.532**	**0.542**	**0.550**	**0.594**	**0.554**

Fig. 7. 3D shape beautification w/o *vs.* w/. our aesthetics scoring. Given input shapes (blue), we display the ablated results of both without (yellow) and with (green) the aesthetic scoring asides.

From Table 1, we observe that our method achieves noticeably higher accuracy than the approach of Dev and Lau's [17] across all four test categories, with more than 10 percentage improvements. We owe this performance gain to the adoption of the ViT backbone in our neural aesthetic model.

6.1 Usability Study

To assess the usability of *Aesthetic3D*, we further conducted two user studies: an ablation study (Section 6.3), an open-ended study (Section 6.4), and two generalization evaluations (Section 6.5).

6.2 Participants

We recruited 18 participants (13 male and 5 female, aged between 23 and 35) from the university email portal and personal networks. Ten of them (S1-S10) are graduate students major in computer science with no prior 3D modeling training, and the remaining eight are professional experts possessing varying levels of 3D modeling experience, including six 3D digital media designers (D1-D6), one artist (A1), and one engineer (E1) developing 3D products with CAD software.

Before the experiment, all participants were given a 15-minute introduction to the functionalities of the *Aesthetic3D* system and the overall procedure of the studies. Each participant was then given 10 minutes to freely explore the interface and tools to become familiar with the system. After this brief orientation, participants proceeded to complete two studies using *Aesthetic3D*. During the studies, a researcher remained present to answer any questions and provide verbal guidance when needed, ensuring the smooth progress of the experiment. After they finished the tasks, we conducted a shorter interview (approximately 10 minutes) with them, to collect their feedback. Note that we recorded the beautification time, usage frequency of each editing tool, and editing steps used to beautify each 3D shape by every participant, for further in-depth analysis.

6.3 Ablation Study

To validate the importance of *Aesthetic3D*'s aesthetic scoring module, we deploy an ablation study where all participants are asked to beautify the same 3D shapes (Figure 7) twice: first without aesthetics scoring, only relying on their intuitions and experience, then with our aesthetic scoring. Note that our ablation follows a within-subject design where both rounds use the same tools with an initial warm-up, so the performance gain in the second round primarily depends on the presence or absence of the aesthetic scoring module rather than increased tool proficiency.

Task. We first disabled the aesthetic scoring module on *Aesthetic3D* interface, asking participants to freely edit 3D objects without any scoring feedback, and save their results (yellow shapes in Figure 7). Next, we enable the scoring component, letting them adjust the same objects using the provided aesthetic scores for guidance until completing their final edits (green shapes in Figure 7).

Qualitative Results. We found that relying solely on participants' intuition (or experience) without aesthetics scoring often led to minor improvements or even degraded the input shapes. For example, in Figure 7 (c), one participant modified the "U"-shape table leg into "11"-shape legs to enhance visual balance, but this adjustment did not yield the highest aesthetic score. By contrast, with our aesthetic scoring, participants could refine their designs through instant feedback, progressively achieving more visually appealing outcomes. Using aesthetics scoring, all participants successfully beautified the input shapes to more aesthetic versions (green shapes in Figure 7).

Regarding time cost, the average editing duration per shape without aesthetics scoring was approximately 4 minutes, whereas those with aesthetics scoring increased to 6 minutes. This additional time reflects participants' exploration of broader beautification directions and an expanded design space enabled by real-time feedback. A portion of the additional time was spent iteratively refining shapes in response to aesthetic guidance. Notably, participants with prior Blender experience (e.g., D4 and D5) completed the beautification task with our aesthetics scoring in just 3 minutes.

In our user interviews, participants expressed strong appreciation for the usefulness of the aesthetic scoring feature. For example, S10 reported that the aesthetic score module was highly helpful for guiding adjustments in shape, proportion, and quantity. D3 also emphasized that the aesthetics scoring effectively guided users toward more aesthetically pleasing direction.

6.4 Open-ended Study

To further explore the potential of *Aesthetic3D* and allow participants greater freedom for creative expression, we conducted an open-ended study.

Task. In this study, participants were asked to select and beautify a set of basic 3D models (e.g., chairs, tables, lamps, mugs) with the full functionality of *Aesthetic3D* system in a freeform manner.

We further evaluate the usability and user experience of *Aesthetic3D* using three questionnaires: (1) a System Usability Scale (SUS) with a 5-point Likert scale (1=strongly disagree, 5=strongly agree); (2) a NASA Task Load Index (NASA-TLX) with a 5-point scale, where lower scores indicate lighter mental and physical workload; and (3) a 5-point Likert scale questionnaire [3] assessing participants' preferences for each aesthetics-driven editing tool and overall evaluation of the system.

Qualitative Results. In this study, participants created a total of 34 beautified results. On average, beautification using *Aesthetic3D* takes approximately 5 minutes, which significantly reduces the time compared to Chen et al. [13] (over 10 minutes). From Figure 8, we observed significant modification or creative beautification was performed on the input shapes, compared to the ablation study (Figure 7). For example, in A1's beautified mug and chair (Figure 8 (b) and (f)), although the original shapes already scored well, A1 can still further enhance them with fancy effects, such as "flower" pattern on the mug and "cat ear" on the chair back. Similarly, D2's beautified results exhibit significant changes, including a smaller lampshade and a more curved lampshade (Figure 8 (e) and (i)).

Figure 8 demonstrates that *Aesthetic3D* effectively supports the creative and aesthetic enhancement across various categories, demonstrating its practical value for different people. Figure 9 depicts participants' editing progression and corresponding aesthetic scores. Each point in the

Fig. 8. Open-ended study: Exhibition of beautified results. For each pair, the left and right represent the shapes before and after beautification. The aesthetic scores and participants' codes are displayed for reference.

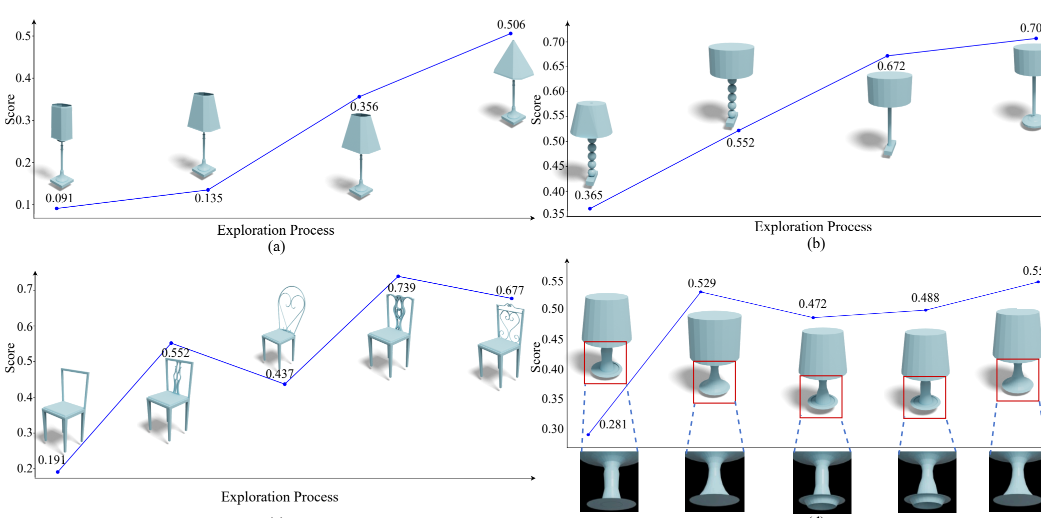

Fig. 9. Open-ended study: four figures showing the results of the participants' explorations.

charts represents a distinct modification, reflecting how design decisions affected aesthetic outcomes during the iterative explorations. For example, in Figure 9 (a), adjusting the lampshade's upper edge led to score variations that guided participants toward a more elegant design.

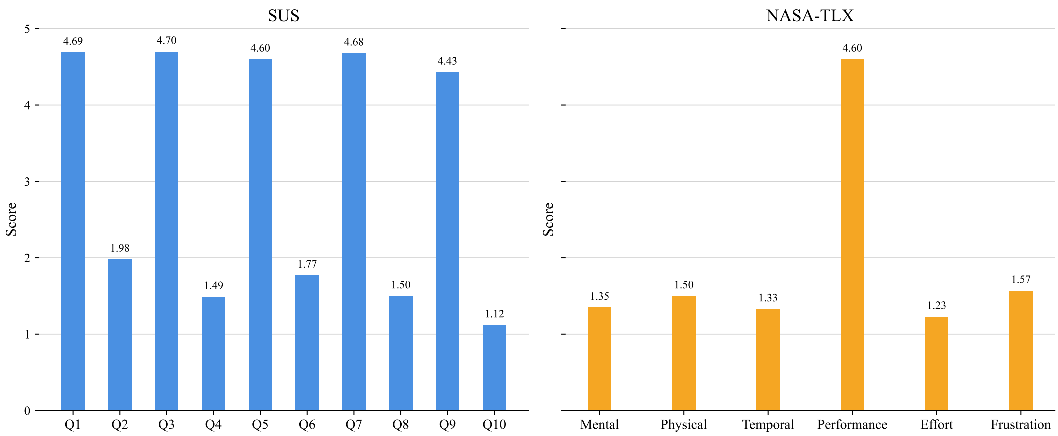

Fig. 10. Mean scores of SUS (left). For the odd-numbered questions, higher scores indicate better performance; for the remaining questions, lower scores indicate better performance. Mean scores of NASA-TLX (right).

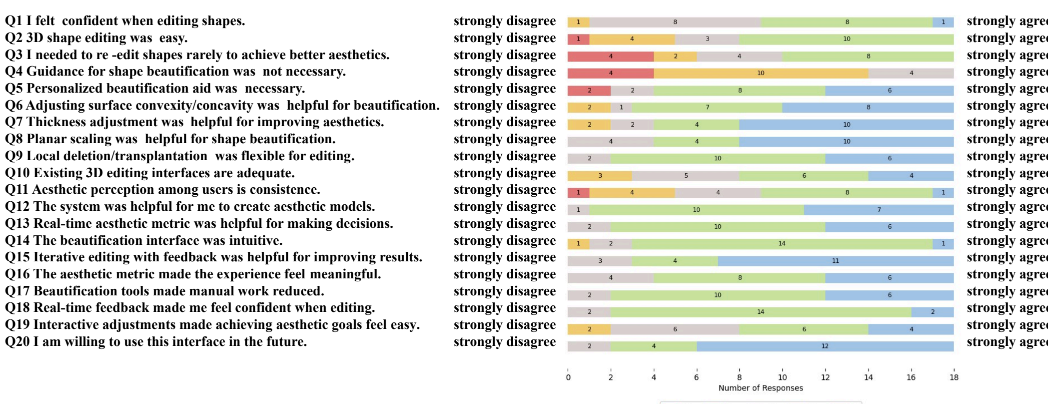

Fig. 11. Open-ended study: questions and participants' voting in the 5-point Likert-scale questionnaire (1–5), describing preferences for individual aesthetics-driven editing tools and overall evaluation of the system.

Quantitative Results. We demonstrate a comprehensive evaluation of *Aesthetic3D* using the SUS and NASA-LTX scores shown in Figure 10. The SUS scores (Figure 10 (left)) indicate that participants generally perceived *Aesthetic3D*'s usability and learnability, reflecting a favorable overall experience. Similarly, the NASA-TLX results (Figure 10 (right)) further suggest relatively low levels of mental, physical, temporal, effort, and frustration.

To report the usefulness of each aesthetics-driven editing tool in *Aesthetic3D*, we analyze responses from the 5-point Likert scale questionnaire (Figure 11). Figure 12 further present four key aspects: participants' preference, usage frequency, usage time for each tool.

(1) Participants' preference. In Figure 12 (left), participants generally held a positive attitude toward the system's tools, with most received average ratings between 3.8 and 4.6. The "*Part Transplant*" tool achieved the highest rating at 4.6, followed by "*Planar Uniform Scaling*" and "*Thickness Adjustment*" (both 4.4). These statistics suggest a preference for the tools producing structural-level modifications.

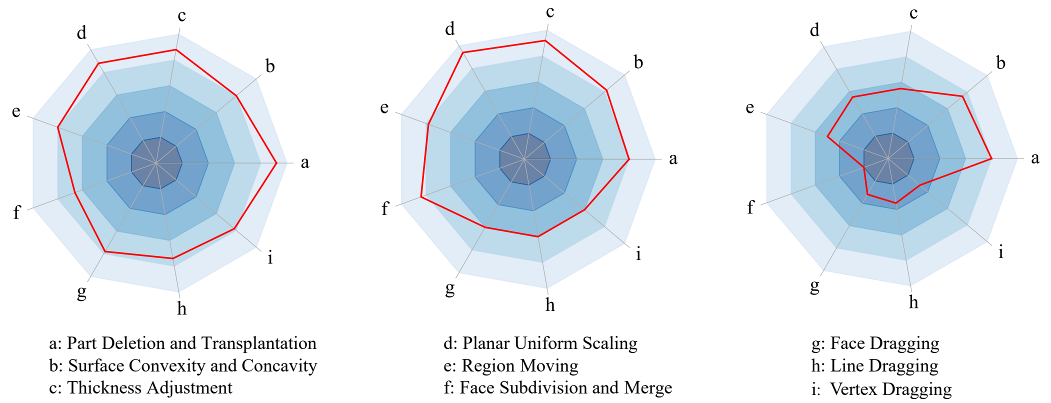

Fig. 12. Three radar charts: distribution of preference ratings for each tool (left), usage frequency for each tool (middle), and time consumption for each tool (right).

(2) Usage frequency. As shown in Figure 12 (middle), tools such as "*Planar Uniform Scaling*", "*Thickness Adjustment*", and "*Face Subdivision and Merge*" were used most frequently, reflecting their widely usage in 3D modeling scenarios. In contrast, finer operations such as "*Vertex Dragging*" and "*Line Dragging*" were used less often, mainly serving as auxiliary functions.

(3) Usage time. From Figure 12 (right), basic tools such as "*Vertex Dragging*", "*Line Dragging*", and "*Face Dragging*" required less than 30 seconds per use, while the more complex tools such as "*Part Transplantation*" and "*Surface Convexity and Concavity Adjustment*" took 60–80 seconds. The statistics suggest that participants tended to rely on basic tools for quick adjustments but invested more time in advanced tools for defining key structural features.

(4) Used Edit Steps. Based on our recordings of editing steps per shape performed by participants, we observe a unimodal distribution, with a peak between 15 and 18 steps (over 20 users), indicating that most participants completed shape beautification within this range. This suggests that *Aesthetic3D*'s editing tools are efficient and allow users to achieve desired results with a reasonable number of steps.

6.5 Generalization Evaluation

To verify the generalization of *Aesthetic3D* beyond the settings of previous user studies, we conducted three generalization evaluations: cross-dataset validation, cross-category validation, and aesthetic tastes on artworks.

Cross-dataset evaluation. As shown in Figure 13, we evaluate *Aesthetic3D* on unseen shapes from the same categories in the ModelNet dataset [50]. Despite noticeable variations in style (Figure 13 (a), (f)) and geometry detail (Figure 13 (c), (d)), the beautified results demonstrate that *Aesthetic3D* can still effectively support aesthetic enhancement even on previously unseen shapes.

Cross-category evaluation. To explore whether our learned aesthetic measures could generalize to other different categories in ModelNet, such as vase, bench, and stool. The beautified results (Figure 14) demonstrate not only the generalization of our learned neural aesthetic measures but also a shared consensus in human perception of aesthetics and beauty across similar categories.

Aesthetic tastes on artworks. To explore how *Aesthetic3D* behaves on shapes that are widely regarded as aesthetic, we apply it to a violin, an art statue, and a product logo, as shown in Figure 15. From the aesthetic scores before and after beautification, we observe that *Aesthetic3D* clearly

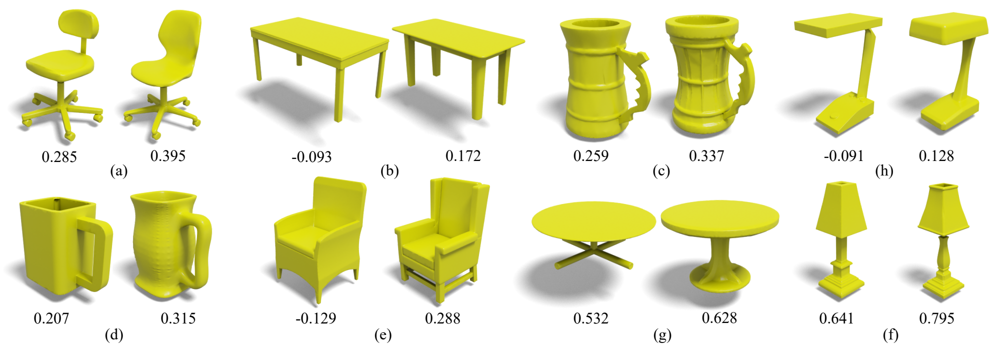

Fig. 13. Cross-dataset validation (ShapeNet→ModelNet). Each pair shows the shapes before and after beautification with *Aesthetic3D*.

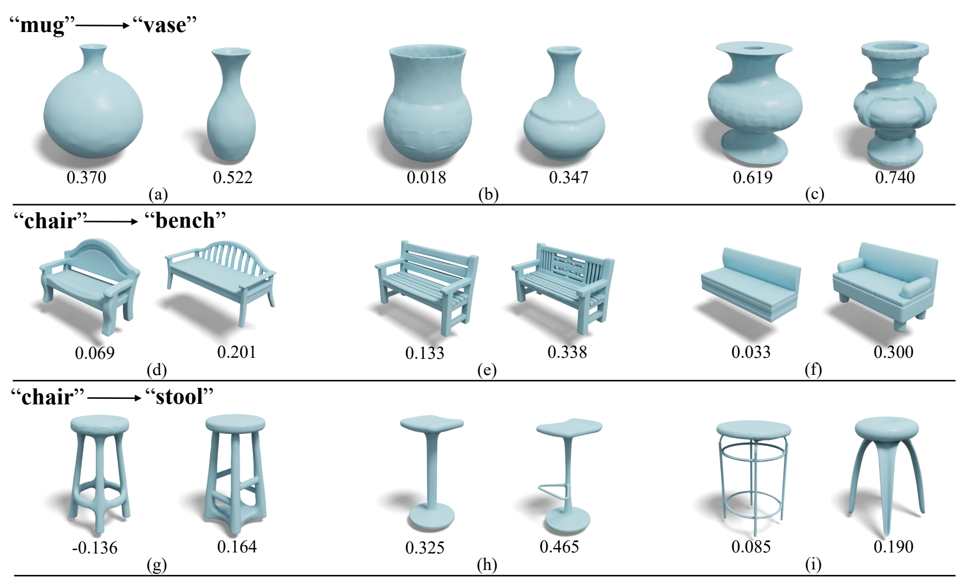

Fig. 14. Cross-category validation. Each pair shows the shapes before and after beautification with *Aesthetic3D*.

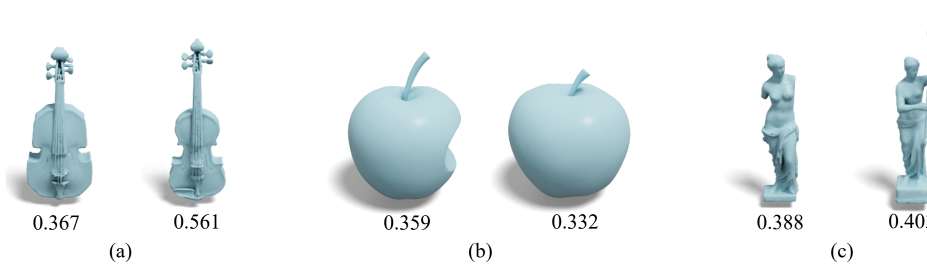

Fig. 15. *Aesthetic3D* applied to artworks that already conform to common aesthetic preferences: (a) violin, (b) Apple logo, and (c) Venus statue, shown before and after beautification.

prefers intact shapes with harmonious structure (Figure 15 (a) and (c)) and favors designs with more curved lines, even when the shape is not fully intact (Figure 15 (b)).

7 Conclusion

In this work, we propose *Aesthetic3D*, a novel 3D shape editing interface that integrates a neural aesthetic measure to enable real-time aesthetics-guided 3D modeling. *Aesthetic3D* quantifies aesthetic features of 3D shapes using a neural network model and incorporates a suite of aesthetics-guided editing tools, allowing users to receive instant feedback and intuitively beautify 3D shapes during modeling. Insights gained from the interviews were useful in the design of *Aesthetic3D*, and could inspire subsequent works for better understanding and studying shape aesthetics further.

While *Aesthetic3D* provides an efficient aesthetic enhancement approach for 3D modeling, our method has some limitations. First, our aesthetic scoring module data-driven and trained on a specific user study dataset, which may constrain its generalization to novel classes with significant shape divergence. This could be partially solved by constructing a larger 3D aesthetics dataset by lifting the 2D image-domain aesthetics dataset to the 3D domain with the help of the Large Language Models (LLMs). Second, our aesthetics-guided editing tools focus more on local geometric modifications, such as convexity adjustments, thickness adjustments, and proportional scaling. In future work, we could explore more complex topological manipulations. Third, our method focuses on capturing common aesthetic consensus for 3D shape. In practice, however, users might want to have a personal aesthetics feedback and even prefer unconventional or polarizing designs. Therefore, our *Aesthetic3D* does not explicitly model bimodal or highly divergent ratings and favors broadly "safe" shapes. One possible solution is to introduce an adaptive mechanism that tailors the common aesthetic measure to individual user preferences. Fourth, the formulation of our *Aesthetic3D* isolates aesthetics from external factors such as material [23, 46], lighting [19], or spatial composition [29], ensuring geometry-focused modeling workflows. Therefore, *Aesthetic3D* will be less effective in scene-context aesthetic scenarios, exploring how an object's appeal emerges within a broader environment, such as composition styles and overall color harmony [31, 48, 49]. Lastly, the interface of our *Aesthetic3D* adopts a straightforward integration of some high-level tools (e.g., convexity/concavity or thickness operations) and low-level operators (e.g., face splitting/merging and vertex dragging). But a more ideal aesthetics-oriented 3D modeling toolset should accommodate diverse aesthetic goals and user expertise, offering multi-level operations guided by aesthetic feedback to lower the entry barrier for novices. Furthermore, incorporating intelligent recommendations could be more friendly and boost creation efficiency.

8 Acknowledgment

We thank the anonymous reviewers for their suggestions. This work was partially funded by National Natural Science Foundation of China (Project No. 62522210), Shandong Provincial Natural Science Foundation (Project No. ZR2025QC1568), Manfred Lau's start-up funds from Lakehead University, Canada, Deng Yu's start-up grant from Shandong University, and Hui Ye's start-up grant from Hong Kong Baptist University.

References

[1] Inc. Autodesk. 2025. Autodesk Maya. https://www.autodesk.com/products/maya. Version 2025, 3D animation and modeling software.

[2] Tunç Ozan Aydın, Aljoscha Smolic, and Markus Gross. 2014. Automated aesthetic analysis of photographic images. *IEEE transactions on visualization and computer graphics* 21, 1 (2014), 31–42.

[3] Seok-Hyung Bae, Ravin Balakrishnan, and Karan Singh. 2009. EverybodyLovesSketch: 3D sketching for a broader audience. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*. 59–68.

[4] Steve Bergen and Brian J Ross. 2013. Aesthetic 3D model evolution. *Genetic Programming and Evolvable Machines* 14, 3 (2013), 339–367.
[5] Francesca Bertacchini, PS Pantano, and Eleonora Bilotta. 2022. Shaping the aesthetical landscape by using image statistics measures. *Acta Psychologica* 224 (2022), 103530.
[6] Dustin Beyer, Serafima Gurevich, Stefanie Mueller, Hsiang-Ting Chen, and Patrick Baudisch. 2015. Platener: Low-fidelity fabrication of 3D objects by substituting 3D print with laser-cut plates. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1799–1806.
[7] Alex Burnap, John R Hauser, and Artem Timoshenko. 2023. Product aesthetic design: A machine learning augmentation. *Marketing Science* 42, 6 (2023), 1029–1056.
[8] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015).
[9] Kathy Charmaz. 2008. Constructionism and the grounded theory method. *Handbook of constructionist research* 1, 1 (2008), 397–412.
[10] Kathy Charmaz. 2014. Constructing grounded theory. (2014).
[11] Siddhartha Chaudhuri, Evangelos Kalogerakis, Leonidas Guibas, and Vladlen Koltun. 2011. Probabilistic reasoning for assembly-based 3D modeling. In *ACM SIGGRAPH 2011 papers*. 1–10.
[12] Minchan Chen and Manfred Lau. 2022. Learning 3d shape aesthetics globally and locally. In *Computer graphics forum*, Vol. 41. Wiley Online Library, 579–588.
[13] Minchan Chen and Manfred Lau. 2024. Enhancing the Aesthetics of 3D Shapes via Reference-based Editing. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–15.
[14] Jiin Choi, Seung Won Lee, and Kyung Hoon Hyun. 2025. GenPara: Enhancing the 3D Design Editing Process by Inferring Users' Regions of Interest with Text-Conditional Shape Parameters. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–21.
[15] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. 2006. Studying aesthetics in photographic images using a computational approach. In *European conference on computer vision*. Springer, 288–301.
[16] Yubin Deng, Chen Change Loy, and Xiaoou Tang. 2018. Aesthetic-driven image enhancement by adversarial learning. In *Proceedings of the 26th ACM international conference on Multimedia*. 870–878.
[17] Kapil Dev and Manfred Lau. 2022. Learning Perceptual Aesthetics of 3-D Shapes From Multiple Views. *IEEE computer graphics and applications* 42, 1 (2022), 20–31.
[18] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR* (2021).
[19] Jarosław Fastowicz, Mateusz Tecław, and Krzysztof Okarma. 2025. A Multi-Feature Automatic Evaluation of the Aesthetics of 3D Printed Surfaces. *Applied Sciences* 15, 9 (2025), 4852.
[20] Guanjun Guo, Hanzi Wang, Chunhua Shen, Yan Yan, and Hong-Yuan Mark Liao. 2018. Automatic image cropping for visual aesthetic enhancement using deep neural networks and cascaded regression. *IEEE Transactions on Multimedia* 20, 8 (2018), 2073–2085.
[21] Ayaan Haque, Matthew Tancik, Alexei A Efros, Aleksander Holynski, and Angjoo Kanazawa. 2023. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF international conference on computer vision*. 19740–19750.
[22] Md Baharul Islam, Wong Lai-Kuan, and Wong Chee-Onn. 2017. A survey of aesthetics-driven image recomposition. *Multimedia Tools and Applications* 76, 7 (2017), 9517–9542.
[23] Yaohui Ji, Yao Sun, Song Qiu, and Xiaojie Geng. 2025. Simultaneous multimodal measures for aesthetic evaluation of furniture color and form. *Scientific Reports* 15, 1 (2025), 2920.
[24] Evangelos Kalogerakis, Siddhartha Chaudhuri, Daphne Koller, and Vladlen Koltun. 2012. A probabilistic model for component-based shape synthesis. *Acm Transactions on Graphics (TOG)* 31, 4 (2012), 1–11.
[25] Vladislav Kreavoy, Dan Julius, and Alla Sheffer. 2007. Model composition from interchangeable components. In *15th Pacific Conference on Computer Graphics and Applications (PG'07)*. IEEE, 129–138.
[26] Congcong Li, Alexander C Loui, and Tsuhan Chen. 2010. Towards aesthetics: A photo quality assessment and photo selection system. In *Proceedings of the 18th ACM international conference on Multimedia*. 827–830.
[27] Dong Liu, Rohit Puri, Nagendra Kamath, and Subhabrata Bhattacharya. 2020. Composition-aware image aesthetics assessment. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 3569–3578.
[28] Ligang Liu, Renjie Chen, Lior Wolf, and Daniel Cohen-Or. 2010. Optimizing photo composition. In *Computer graphics forum*, Vol. 29. Wiley Online Library, 469–478.
[29] Tianqiang Liu, Aaron Hertzmann, Wilmot Li, and Thomas Funkhouser. 2015. Style compatibility for 3D furniture models. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 1–9.

[30] Zhen Liu, Yao Feng, Michael J Black, Derek Nowrouzezahrai, Liam Paull, and Weiyang Liu. 2023. Meshdiffusion: Score-based generative 3d mesh modeling. *arXiv preprint arXiv:2303.08133* (2023).

[31] Zhaoliang Lun, Evangelos Kalogerakis, and Alla Sheffer. 2015. Elements of style: learning perceptual shape style similarity. *ACM Transactions on graphics (TOG)* 34, 4 (2015), 1–14.

[32] Mayra D Barrera Machuca, Paul Asente, Wolfgang Stuerzlinger, Jingwan Lu, and Byungmoon Kim. 2018. Multiplanes: Assisted freehand vr sketching. In *Proceedings of the 2018 ACM Symposium on Spatial User Interaction*. 36–47.

[33] Ian Mason, Sebastian Starke, and Taku Komura. 2022. Real-time style modelling of human locomotion via feature-wise transformations and local motion phases. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 5, 1 (2022), 1–18.

[34] Maxon. 2025. Cinema 4D. https://www.maxon.net/cinema-4d. 3D modeling, animation, and rendering software.

[35] Maxon. 2025. ZBrush. https://www.maxon.net/en/zbrush. Digital sculpting and painting software.

[36] Jon McCormack and Andy Lomas. 2021. Deep learning of individual aesthetics. *Neural Computing and Applications* 33, 1 (2021), 3–17.

[37] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 4460–4470.

[38] Kenjiro T Miura and Gobithaasan RU. 2014. Aesthetic curves and surfaces in computer aided geometric design. *International Journal of Automation Technology* 8, 3 (2014), 304–316.

[39] Masashi Nishiyama, Takahiro Okabe, Imari Sato, and Yoichi Sato. 2011. Aesthetic quality classification of photographs based on color harmony. In *CVPR 2011*. IEEE, 33–40.

[40] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2015. Designscape: Design with interactive layout suggestions. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1221–1224.

[41] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. 2019. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 165–174.

[42] Binh Pham. 1999. Design for aesthetics: interactions of design variables and aesthetic properties. In *Human vision and electronic imaging IV*, Vol. 3644. SPIE, 364–371.

[43] Binh Pham and Jinglan Zhang. 2003. A fuzzy shape specification system to support design for aesthetics. In *Soft Computing in Measurement and Information Acquisition*. Springer, 39–50.

[44] Carlo H Séquin. 2005. CAD tools for aesthetic engineering. *Computer-aided design* 37, 7 (2005), 737–750.

[45] Zekun Sun and Chaz Firestone. 2022. Beautiful on the inside: Aesthetic preferences and the skeletal complexity of shapes. *Perception* 51, 12 (2022), 904–918.

[46] Jieni Tian, Jiangping Yuan, Hua Li, Danyang Yao, and Guangxue Chen. 2021. Advanced surface color quality assessment in paper-based full-color 3D printing. *Materials* 14, 4 (2021), 736.

[47] Cesar Torres, Tim Campbell, Neil Kumar, and Eric Paulos. 2015. HapticPrint: Designing feel aesthetics for digital fabrication. In *Proceedings of the 28th annual ACM symposium on user interface software & technology*. 583–591.

[48] Shun-Yu Wang, Wei-Chung Su, Serena Chen, Ching-Yi Tsai, Marta Misztal, Katherine M Cheng, Alwena Lin, Yu Chen, and Mike Y Chen. 2024. Roomdreaming: Generative-AI approach to facilitating iterative, preliminary interior design exploration. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.

[49] Tomer Weiss, Ilkay Yildiz, Nitin Agarwal, Esra Ataer-Cansizoglu, and Jae-Woo Choi. 2020. Image-driven furniture style for interactive 3D scene modeling. In *Computer Graphics Forum*, Vol. 39. Wiley Online Library, 57–68.

[50] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 2015. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[51] Kai Xu, Hao Zhang, Daniel Cohen-Or, and Baoquan Chen. 2012. Fit and diverse: Set evolution for inspiring 3d shape galleries. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 1–10.

[52] Pengfei Xu, Guohang Yan, Hongbo Fu, Takeo Igarashi, Chiew-Lan Tai, and Hui Huang. 2019. Global beautification of 2D and 3D layouts with interactive ambiguity resolution. *IEEE transactions on visualization and computer graphics* 27, 4 (2019), 2355–2368.

[53] Zhicheng Yan, Hao Zhang, Baoyuan Wang, Sylvain Paris, and Yizhou Yu. 2016. Automatic photo adjustment using deep neural networks. *ACM Transactions on Graphics (TOG)* 35, 2 (2016), 1–15.

[54] Yuzhe Yang, Liwu Xu, Leida Li, Nan Qie, Yaqian Li, Peng Zhang, and Yandong Guo. 2022. Personalized image aesthetics assessment with rich attributes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 19861–19869.

[55] Kangxue Yin, Jun Gao, Maria Shugrina, Sameh Khamis, and Sanja Fidler. 2021. 3dstylenet: Creating 3d shapes with geometric and texture style variations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12456–12465.

[56] Deng Yu, Manfred Lau, Lin Gao, and Hongbo Fu. 2024. Sketch Beautification: Learning Part Beautification and Structure Refinement for Sketches of Man-Made Objects. *IEEE Transactions on Visualization and Computer Graphics* 30, 9 (2024), 6533–6546.

[57] Deng Yu, Chufeng Xiao, Manfred Lau, and Hongbo Fu. 2024. Sketch2Stress: Sketching With Structural Stress Awareness. *IEEE Transactions on Visualization and Computer Graphics* 30, 10 (2024), 6851–6865.

[58] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. 2022. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 18353–18364.

[59] Jie Zhou, Linzi Qu, Miu-Ling Lam, and Hongbo Fu. 2025. From Rigging to Waving: 3D-Guided Diffusion for Natural Animation of Hand-Drawn Characters. *ACM Transactions on Graphics (TOG)* 44, 6 (2025), 1–11.

[60] Shizhe Zhou, Hongbo Fu, Ligang Liu, Daniel Cohen-Or, and Xiaoguang Han. 2010. Parametric Reshaping of Human Bodies in Images. *ACM Transactions on Computer Graphics: Special Issue of ACM SIGGRAPH 2010* 29, 4 (2010).

Supplementary Materials: *Aesthetic3D*: Incorporating Shape Aesthetic Measures into 3D Modeling Interfaces

DENG YU and JIANING GUO, Shandong University
HUI YE, Hong Kong Baptist University
PENGJIE REN*, Shandong University and Qingdao Key Laboratory of Trustworthy Artificial Intelligence
HONGBO FU*, Hong Kong University of Science and Technology
MANFRED LAU*, Lakehead University

ACM Reference Format:
Deng Yu, Jianing Guo, Hui Ye, Pengjie Ren, Hongbo Fu, and Manfred Lau. 2026. Supplementary Materials: *Aesthetic3D*: Incorporating Shape Aesthetic Measures into 3D Modeling Interfaces. *Proc. ACM Comput. Graph. Interact. Tech.* 9, 1, Article 9 (May 2026), 4 pages. https://doi.org/10.1145/3804501

1 Formative Study

1.1 Interview Questions

(1) Background and Objectives (Introduction)

Purpose: To understand participants' backgrounds and their relevance to the research topic.

- Could you briefly introduce your role, work, or experience? How many years have you been working in this field?
- Have you ever been involved in **aesthetic design of 3D shapes**? What kind of experience do you have?
- What are your main goals or expectations regarding **aesthetic enhancement technologies** for 3D objects?

(2) Current Behaviors and Needs

Purpose: To explore users' existing practices, pain points, and unmet needs.

- How do you judge whether **a 3D object is aesthetically pleasing**? Do you encounter situations where team members or supervisors have inconsistent standards?
- How do you **currently beautify or enhance a 3D object**? What tools do you use?
- What is the biggest challenge or frustration you experience in this process?
- Are there any existing solutions? What are their limitations?
- How would you like to improve your current workflow or tools?

*Corresponding authors

Authors' Contact Information: Deng Yu, deng.yu@sdu.edu.cn; Jianing Guo, Shandong University; Hui Ye, huiyehy@hkbu.edu.hk, Hong Kong Baptist University; Pengjie Ren, Shandong University and Qingdao Key Laboratory of Trustworthy Artificial Intelligence; Hongbo Fu, fuplus@gmail.com, Hong Kong University of Science and Technology; Manfred Lau, manfred.lau@lakeheadu.ca, Lakehead University.

ACM 2577-6193/2026/5-ART9
https://doi.org/10.1145/3804501

(3) Attitudes and Perceptions
Purpose: To understand participants' perspectives and potential biases regarding the research topic.

- What is your first impression of **3D aesthetic enhancement**?
- What do you think is the core value of **aesthetic scoring**?
- Are there other factors (e.g., cultural background, resource constraints) that influence your view of **aesthetic design**?

(4) Specific Scenarios and Use Cases
Purpose: To gather detailed insights through concrete scenarios.

- Could you describe a recent case where you used an **aesthetic enhancement tool**?
- In this case, which aspects did you find efficient or inefficient?
- If you could have an ideal **interactive aesthetic enhancement tool**, what features would you expect it to have?

(5) Feedback and Iteration Purpose: To collect feedback and improvement suggestions on existing prototypes or concepts (if available).

- (Show prototype/concept) Does this design meet your expectations? Why or why not?
- Which aspects do you think need the most improvement?
- Are there any additional features you would like to see included?

(6) Closing and Additional Input
Purpose: To ensure no important topics are missed.

- Are there any other issues we haven't discussed that you think are important?
- Are there other people or roles you would suggest we interview?

2 Method

2.1 Details of our Aesthetics-driven Editing Tools

We detailed the instructions for the usage of our 9 types of aesthetics-driven editing tools below:

(a) *Surface Convexity and Concavity Adjustment*: supports precise modification of surface curvature by adjusting convexity and concavity, or sculpting relatively flat regions to enhance the hierarchical richness of the form.

(b) *Region Thickness Adjustment*: provides localized control over the thickness of structural parts, facilitating proportional refinement of vertical or elongated components (e.g., chair legs, table posts).

(c) *Region Moving*: allows the selection of connected regions composed of vertices, edges, and faces, which can then be translated to different positions within the model, supporting mid-level structural rearrangements.

(d-e) *Face Subdivision or Merge*: provides dual capabilities for surface granularity control. Subdivision splits larger faces into smaller ones to create finer editing units, while merging integrates adjacent planar faces into larger regions for holistic editing.

(f) *Planar Uniform Scaling*: allows for uniform scaling of planar elements such as seats, backs, or tabletops, enabling users to adjust relative proportions of planar versus vertical components in a controlled manner.

(g) *Vertex Dragging*: provides the most fine-grained level of control, enabling users to directly move vertices for detailed adjustments of geometry.

(h) *Line Dragging*: allows users to reshape geometry along edges, enabling precise adjustments of linear structural features.

(i) *Face Dragging*: enables the direct manipulation of surface elements by translating selected faces to adjust the geometry at the surface level.

(j) *Part Transplant*: enables users to remove redundant local parts or transplant substructures from one region to another, thereby reshaping the overall composition while maintaining aesthetic coherence.

3 Evaluation

3.1 Questionnaires

Here, we elaborate on the questions used in the three questionnaires with a 5-point scale, namely, the system usability scale (SUS) questionnaire, the NASA Task Load Index (NASA-TLX) questionnaire, and the questionnaire specifically designed to evaluate our *Aesthetic3D* system and aesthetics-driven editing tools, respectively. Note that 5-point scale answers are omitted for simplicity, only describing the questions below.

For the SUS questionnaire, we invite participants to rate the questions below:
Q1: "*I think that I would like to use this system frequently.*"
Q2: "*I found the system unnecessarily complex.*"
Q3: "*I thought the system was easy to use.*"
Q4: "*I think that I would need the support of a technical person to be able to use this system.*"
Q5: "*I found the various functions in this system were well integrated.*"
Q6: "*I thought there was too much inconsistency in this system.*"
Q7: "*I would imagine that most people would learn to use this system very quickly.*"
Q8: "*I found the system very cumbersome to use.*"
Q9: "*I felt very confident using the system.*"
Q10: "*I needed to learn a lot of things before I could get going with this system.*"

Regarding the NASA-TLX questionnaire, we invite participants to rate the questions below:
Q1: "*Mental Demand: How much mental and perceptual activity was required* ?"
Q2: "*Physical Demand: How much physical activity was required*?"
Q3: "*Temporal Demand: How hurried or rushed was the pace of the task*?"
Q4: "*Performance: How successful were you in accomplishing what you were asked to do*?"
Q5: "*Effort: How hard did you have to work to accomplish your level of performance*?"
Q6: "*Frustration: How insecure, discouraged, irritated, stressed, and annoyed were you*?"

As for the questions for evaluating our *Aesthetic3D* and aesthetics-driven editing tools, we describe them below:
Q1: "*I felt confident when editing shapes.*"
Q2: "*3D shape editing was easy.*"
Q3: "*I needed to re-edit shapes rarely to achieve better aesthetics.*"
Q4: "*Guidance for shape beautification was not necessary.*"
Q5: "*Personalized beautification aid was necessary.*"
Q6: "*Adjusting surface convexity/concavity was helpful for beautification.*"
Q7: "*Thickness adjustment was helpful for improving aesthetics.*"
Q8: "*Planar uniform scaling was helpful for shape beautification.*"
Q9: "*Part transplant was flexible for editing.*"
Q10: "*Existing 3D editing interfaces are adequate.*"
Q11: "*Aesthetic scoring among users is consistence.*"
Q12: "*The system was helpful for me to create aesthetic models.*"

Q13: "*Real-time aesthetic scoring was helpful for making decisions.*"
Q14: "*The beautification interface was intuitive.*"
Q15: "*Iterative editing with feedback was helpful for improving results.*"
Q16: "*The aesthetic scoring made the experience feel meaningful.*"
Q17: "*The aesthetics-driven editing tools reduced manual workload.*"
Q18: "*Real-time feedback made me feel confident during 3D modeling.*"
Q19: "*Interactive adjustments made achieving aesthetic goals easy.*"
Q20: "*I am willing to use this interface in the future.*"