

DreamCAD: Scaling Multimodal CAD Generation using Differentiable Parametric Surfaces

Anonymous ECCV 2026 Submission

Paper ID #3740

Abstract. Multimodal CAD generation faces a fundamental scalability challenge. Design-history methods are confined to small annotated datasets, while BRep topology is discrete and non-differentiable. Meanwhile, millions of unannotated 3D meshes remain untapped, since existing CAD methods cannot leverage them without explicit CAD annotations. We propose **DreamCAD**, a multimodal generative framework that bridges this gap by representing shapes as C^0 -continuous Bézier patches with differentiable tessellation, enabling direct point-level supervision on large-scale 3D meshes without CAD-specific annotations. The resulting surfaces are exportable as STEP files and editable in standard CAD software. We further introduce **CADCap-1M**, the largest CAD captioning dataset with 1M+ GPT-5-generated descriptions to advance text-to-CAD research. DreamCAD achieves state-of-the-art performance on ABC and Objaverse across text, image, and point modalities, surpassing 75% user preference. Finally, we demonstrate that DreamCAD’s accurate and compact geometric foundation enables topology recovery for production-ready CAD generation.

Keywords: Multimodal CAD · Parametric Surfaces · Caption Dataset

1 Introduction

Computer-Aided Design forms the foundation of modern engineering [73], architecture, and manufacturing [28]. Unlike 3D meshes or point clouds that approximate geometry [59], CAD models are built from parametric primitives such as Bézier and NURBS surfaces, encoded in boundary representation (BRep) format, capturing both topology and precise geometry essential for industrial design [6]. With the advent of generative AI [10, 65], AI-assisted CAD generation [39] has become a promising frontier for accelerating design workflows and enabling rapid prototyping of manufacturable models [17].

Recent CAD generation methods enable multimodal synthesis from text [21], images [69], or point clouds [35]. However, achieving strong generalization across modalities and diverse geometries remains a major challenge. The root cause lies in the CAD representation and the training scheme. Design-history-based



Fig. 1: Our proposed **DreamCAD** (left) is a multimodal generative framework that can reconstruct CAD models from text, images, and point clouds using parametric patches. **CADCap-1M** (right) provides 1M+ GPT-5-generated captions.

models [62] rely on *sketch-and-extrude* sequences from small datasets such as DeepCAD-160k [62] and Fusion360-8k [60], constraining the generalization in freeform or open-vocabulary shapes. UV and graph-based methods [67] require explicit BRep annotations, which are costly and difficult to scale. Although ABC [23] contains 1M BReps, it remains underutilized due to these bottlenecks.

In contrast, modern 3D generative models [71] scale effectively by decoupling mesh generation into multiple stages, using intermediate representations such as SDFs [47], or NeRFs [40]. We argue that scaling multimodal CAD generation requires an analogous paradigm shift - from single-stage joint geometry and topology generation to a decoupled pipeline. The first stage should focus on scalable parametric geometry leveraging unstructured 3D data without explicit CAD annotations, while subsequent stages can work toward recovering full CAD topology from this intermediate geometric foundation.

To address the first stage of generalizability, we propose **DreamCAD**, a multimodal generative framework that produces editable parametric surfaces from point-level supervision alone. Each CAD model is represented as a set of rational Bézier patches defined by learnable control points and weights. These surfaces are differentially tessellated into meshes for point-based supervision via Chamfer loss. DreamCAD encodes sparse voxels into structured latents (SLAT) [65] to learn a CAD-oriented latent space, which is decoded into parametric geometry. Building on this, we adopt a coarse-to-fine conditional generation framework supporting *text*, *image*, and *point cloud* inputs.

A key challenge with patch-based BRep representations is ensuring C^0 continuity, which is essential for valid CAD modeling. C^0 continuity [11] requires adjacent surface patches to share common boundaries without gaps or overlaps. We address the continuity problem structurally rather than using any geometric optimization. Starting from a sparse voxel grid, we remove internal quads

Table 1: Training datasets for DreamCAD comprising 1.3M 3D meshes. ✓ denotes if the dataset is included in proposed CADCap-1M dataset.

Dataset	Samples	CADCap	Dataset	Samples	CADCap
ABC [23]	757,433	✓	3D-Future [13]	16,990	✓
Automate [19]	380,124	✓	ModelNet [64]	12,308	✓
ShapeNet [4]	52,458	✗	ABO [5]	7,944	✗
CADParser [72]	40,989	✓	Fusion360 [60]	4,603	✓
HSSD [22]	30,078	✗	Toys4K [54]	3,482	✗
Total			1,306,409		

via flood-fill [12]. This way, each surface quad maps to a parametric patch initialized with control-point grids and unit weights, with adjacent patches sharing boundary points for continuity. The VAE decoder then refines control points and weights to match the target geometry. At inference, Bézier surfaces are exported as STEP files via OpenCascade [44]. Importantly, while the resulting STEP files do not contain industry-standard CAD topology, they are directly editable via control point and weight manipulation. In Section 5.4, we show that DreamCAD’s high-fidelity and compact geometric reconstruction can provide a strong geometric prior for topology recovery.

We curate over 1M 3D meshes from 10 publicly available datasets, as summarized in Table 1, to train our VAE architecture. For *text-to-CAD* training, existing datasets such as Text2CAD-160K [21] lack the scale and diversity necessary for robust generative modeling. To address this, we construct **CADCap-1M**, the largest CAD captioning dataset to date. We leverage GPT-5 to generate high-quality captions for 1M+ models from existing large-scale CAD datasets. We evaluate DreamCAD on multimodal generation tasks across *text*, *image*, and *point cloud* on both ABC [23] and Objaverse [7] datasets. Our experiments show that DreamCAD consistently surpasses baselines in both geometric accuracy and user preference. Our contributions are as follows

1. We introduce **DreamCAD**, a multi-modal generative framework trained only with point supervision without any dependence on CAD annotations.
2. We release **CADCap-1M**, the largest CAD captioning dataset with over 1M text descriptions for scalable text-to-CAD research.
3. DreamCAD achieves state-of-the-art performance across *point*-, *image*-, and *text*-conditioned generation tasks, reducing Chamfer Distance by up to 70% in *point-to-CAD* and surpassing 75% preference in expert and GPT-based evaluations for *text* and *image-to-CAD*.

2 Related Work

Generative CAD. Prior CAD generative models [26,55,62] represent geometry through boundary representations (BReps), which offer exact, watertight, and editable geometry but are difficult to learn at scale due to their complex parametric and topological structure [15]. Existing approaches therefore struggle to balance scalability with geometric fidelity.

Design-history methods [21, 50, 62] model CAD creation using *sketch-extrude* operations, framed as program synthesis [50, 58, 66, 72] or language modeling [9, 20, 29, 38]. However, they require design-history logs available only in proprietary CAD softwares [2, 43]. This limits training to small datasets (e.g., DeepCAD-170K [62], Fusion360-8K [60]) and models exhibit poor generalization to complex, real-world geometries [14, 58].

Some works model BReps through UV parameterizations or hierarchical graph structures since explicit BRep topology (e.g. vertices, edges, faces) is discrete and non-differentiable, preventing gradient-based optimization for geometry based learning. UVNet [18] and DTG-BRepGen [31] predict UV grids from face annotations, while BRepGen [67] and BRepDiff [27] perform UV-space diffusion with post-processing. However, these methods face some key limitations: (i) no guaranteed C^0 continuity across adjacent faces, (ii) training restricted to models with fewer faces [27], excluding $\sim 70\%$ of ABC [23], (iii) BRep conversion requires expensive grid-based fitting stages with high error and invalidity rate [1, 15], and (iv) UV parameterizations only approximate geometry with resolution-related computational challenges. Self-supervised methods [32, 49, 56] attempts to learn CAD geometry without using any CAD annotation. However, these methods either do not scale well or restrict to simpler shapes and have never been explored for multimodal purposes. Similarly, surface-fitting approaches [35, 51, 61] require slow per-sample optimization, unsuitable for scalable generation. NURBGen [55] takes a promising direction by framing BRep generation as a sequence modeling task using symbolic NURBS, enabling text-to-CAD via LLM fine-tuning. However, generating 3D geometry directly from text without visual grounding leads to low fidelity [52, 53] on complex or geometrically precise prompts.

While industry-standard BRep topology remains the gold standard for professional CAD workflows, directly training on BReps from multimodal inputs is neither scalable nor viable, as the above limitations collectively prevent gradient-based learning on large-scale BRep data. Inspired by recent progress in 3D generation [52, 70], we argue that scalable CAD synthesis requires a decoupled two-stage pipeline: a first stage that learns generalizable 3D shape from large-scale unstructured data, and a second stage that recovers fine-grained CAD topology from these intermediate representations. DreamCAD addresses the first stage via C^0 -continuous Bézier patches with differentiable tessellation, enabling point-level supervision on large-scale 3D meshes without CAD-specific annotations.

Multimodal CAD Datasets. Unlike 3D vision, where large-scale multimodal datasets such as Objaverse [7] and MARVEL-40M+ [53] have driven progress in text-to-3D generation, the CAD domain remains limited by multimodal data scarcity. Existing CAD datasets, such as Text2CAD-160K [21], contain design histories with text captions but are small in scale. Moreover, large-scale CAD datasets such as ABC-1M [23] and Automate-440K [19] lack textual or visual descriptions, restricting their use for multimodal generative learning. Although recent advances in automated 3D captioning, including Cap3D [37] and MARVEL-40M+ [53], have enabled large-scale annotations for meshes, there are still no comparable resources for BRep models. To bridge this gap, we in-




Fig. 2: Bézier surface representation and differentiable tessellation.

troduce **CADCap-1M**, a dataset of over 1M high-quality text descriptions for CAD models automatically generated using GPT-5, for scalable training and evaluation for text-to-CAD research.

3 Preliminaries

In this section, we briefly review the CAD representation used in the DreamCAD architecture. Among the various parametric surface formulations, Bézier and NURBS are the most widely adopted in modern CAD modeling. In our implementation, we choose bicubic rational Bézier surfaces due to their conceptual simplicity and analytical tractability. Furthermore, rational Bézier surfaces can be viewed as a special case of NURBS surfaces, making them naturally compatible with standard CAD operations.

Rational Bézier Surface. A rational Bézier surface of degree (n, m) is defined by $(n+1) \times (m+1)$ control points $C = \{c_{ij}\}$ and non-negative weights $W = \{w_{ij}\}$. The surface $S(u, v)$ is evaluated in uv domain as:

$$S(u, v) = \frac{\sum_{i,j} B_i^n(u) B_j^m(v) w_{ij} c_{ij}}{\sum_{i,j} B_i^n(u) B_j^m(v) w_{ij}}, \quad (1)$$

where $B_i^n(u) = \binom{n}{i} u^i (1-u)^{n-i}$ and $B_j^m(v) = \binom{m}{j} v^j (1-v)^{m-j}$ are Bernstein basis functions and $(u, v) \in [0, 1]^2$. For a bicubic case $n = m = 3$. This formulation is differentiable with respect to both control points and weight. It is worth noting that the weights w_{ij} must remain non-negative, as negative weights can lead to degenerate or invalid surface evaluations.

Differentiable Mesh Generation. Given a set of Bézier patches $\{S_k\}_{k=1}^K$, we generate meshes through differentiable tessellation [61] (Figure 2). Each patch $S_k(u, v)$ is evaluated on the uv domain by uniformly sampling (u, v) points on a grid of resolution $r \times r$. Adjacent points in this grid define quadrilateral cells, which are then split into triangles to form a locally consistent mesh. Neighboring Bézier patches are merged along shared boundaries to ensure C^0 connectivity. Since $S(u, v)$ is differentiable with respect to its C and W , the entire tessellation process supports end-to-end gradient-based optimization.

4 Methodology

As shown in Figure 3, **DreamCAD** adopts a multi-stage generative pipeline. Section 4.1 presents the VAE module, which encodes 3D shapes into compact




Fig. 3: DreamCAD Overview: (A). Sparse Transformer VAE takes as input mesh, generates active voxels v_i with local features f_i , from DINOv2 [46] embeddings, normal images, and SDF values and encodes it to generate structured latents z_i . These are then decoded into parametric (rational bézier) surfaces and optimized using Chamfer loss. **(B). Initial C^0 -continuous Parametric Surface** generation from sparse voxels via flood-fill and quad conversion using grid control points and unit weights. **(C). Multi-modal CAD generation** from images, or points using a coarse-to-fine flow-matching framework from coarse voxel grid to parametric surface refinement.

latent representations and decodes them into a parametric surface. Section 4.2 details the automatic captioning pipeline. Finally, Section 4.3 introduces the conditional generation framework, which uses a coarse-to-fine strategy: sparse voxels are first generated from the input condition, followed by reconstruction of fine-grained parametric surfaces.

4.1 Latent Encoding

Sparse Voxel Representation. An effective 3D encoder requires a compact yet structured representation. While point clouds are widely used [33, 71], they lack spatial regularity for continuous CAD generation. We instead adopt a sparse voxel representation enriched with local visual features [65]. We first voxelize each input mesh to 32^3 resolution, generating active voxels $\{v_i\}_{i=1}^N$. To preserve fine geometric details, each active voxel is augmented with visual cues. We render 150 RGB and normal views of the mesh from different camera angles, extract DINOv2 [46] embeddings, and project each voxel center v_i to obtain per-view RGB and normal features. Averaging across views gives:

$$p(v_i) = \frac{1}{150} \sum_{j=1}^{150} \left[\text{proj}(v_i, E_j^r); \text{proj}(v_i, E_j^n) \right], \quad (2)$$

where proj is the projection operator and $p(v_i)$ is the mean feature vector, E_j^r and E_j^n are DINO embedding maps from the j -th RGB and normal views. We further include per-view normals $n(v_i) \in \mathbb{R}^{150 \times 3}$, voxel centers $c(v_i) \in \mathbb{R}^3$, and signed distance values $s(v_i) \in \mathbb{R}$ to encode geometry and surface proximity. The final voxel feature is:

$$f_i = [p(v_i); n(v_i); c(v_i); s(v_i)], \quad f_i \in \mathbb{R}^{2502} \quad (3)$$

These features are processed by a sparse Transformer encoder [65] to produce structured latents $\{(v_i, z_i)\}_{i=1}^N$.

Initial Parametric Quad Generation. The goal of the decoder is to reconstruct the 3D shape as a set of Bézier patches from the structured latent codes. Each patch requires 16 control points and corresponding weights. Directly generating these parameters from latents z_i will lead to disconnected or overlapping patches, as losses alone cannot enforce C^0 continuity between two adjacent patches.

To mitigate this, we generate an initial parametric quad surface from sparse voxels (Fig. 3.B). First, we generate a surface mesh from the sparse voxels using a flood-fill algorithm, which removes the internal quads from the voxels. We then convert each quad into a bicubic rational bézier patch by uniformly sampling a 4×4 grid points using bilinear interpolation of its four corners.

All control points start with unit weights, and adjacent patches share boundary control points (“● Shared” in Figure 3.A) to ensure C^0 continuity. The resulting surface S is represented as:

$$\mathcal{S} = \{s_k = (c_{ij}^k, w_{ij}^k)\} \in \mathbb{R}^{N_f \times 16 \times 4}, \quad k \in \mathbb{Z}_{[0, N_f]}, \quad i, j \in \{0, 1, 2, 3\}, \quad (4)$$

where N_f denotes the number of patches, s_k denotes the k^{th} patch and i, j are the indices for both control points and weights. In practice, we observe that the number of patches per shape remains low, typically $N_f \ll 10000$.

Parametric Surface Decoder. The decoder refines the initial surface $\mathcal{S}$ by using structured latent features to predict local adjustments for each patch. For the (i, j) -th control point of the k -th patch, it predicts a deformation d_{ij}^k and a weight update $\hat{w}_{ij}^k$. Unconstrained d_{ij}^k often causes degenerate geometries such as spikes or self-intersections, which are hard to recover during optimization. Therefore, to stabilize training, we bound the deformation within a local neighborhood using $c_{ij}^k \leftarrow c_{ij}^k + \tanh(d_{ij}^k)$. We ensure positive weights via the softplus function $\log(1 + e^{\hat{w}_{ij}^k})$. For shared boundary control points, we enforce C^0 continuity by uniformly averaging the predicted deformations and weight updates from all patches sharing that point, ensuring all adjacent patches converge to the same boundary position.

We then tessellate the deformed surface into a mesh and compute Chamfer distance (CD) loss between sampled surface points ($\mathcal{X}_d$) and the target point cloud ($\mathcal{X}_g$). The overall training objective is:

$$\mathcal{L} = \lambda_{cd} \text{CD}(\mathcal{X}_g, \mathcal{X}_d) + \lambda_{g1} \text{G1}(\mathcal{S}_d) + \lambda_{lp} \text{Laplacian}(\mathcal{M}_d) + \lambda_{kl} D_{KL}, \quad (5)$$

where $\mathcal{M}_d$ is the tessellated mesh from $\mathcal{S}_d$, **G1** [8] enforces tangent continuity between two adjacent patches, **Laplacian** [42] ensures smoothness and D_{KL} regularizes the VAE latent space using KL divergence loss.

4.2 CADCap-1M Dataset

Recent progress in text-to-3D generation has been driven by large-scale captioned datasets such as MARVEL-40M+ [53] and Cap3D [37], yet no comparable



Fig. 4: Examples of metadata-augmented captions from CADCap-1M showing object type, part names, and hole counts.

resource exists for *text-to-CAD*. To progress the field, we introduce CADCap-1M, comprising 1M+ high-quality captions for CAD models from ABC [23], Automate [19], CADParser [72], Fusion360 [60], ModelNet [64], 3D-Future [13]. For each model, we render four orthographic views using Blender [3] and prompt GPT-5 [45] to generate concise descriptions. Prompts are augmented with metadata from the original CAD files, such as model names which are optionally extracted from *.step* files, number of holes, which is computed using [41], and relative dimensions (*length-to-width*, *width-to-height*, or *length-to-height* ratios). This metadata-augmented prompting substantially reduces hallucinations and improves geometric accuracy and linguistic quality [21, 53]. As a result, CADCap-1M produces more shape-centric and structure-aware captions (e.g., “*M3x8 bolt ..*”, “*.. mounting plate .. 40 circular holes ..*”), as illustrated in Figure 4.

4.3 Conditional CAD Generation

We now describe the process of conditional CAD generation from text, images, and point clouds.

Generation Pipeline. As shown in Figure 3.C, DreamCAD adopts a coarse-to-fine generation pipeline with two Flow Transformer Decoders optimized via flow matching (FM) objectives [34]. FM learns to estimate velocity field $v_\theta(x_t, t)$ for transforming samples from a prior distribution $x_0 \sim p_0$ to the target distribution $x_1 \sim p_1$ through the loss $\mathcal{L}_{\text{FM}} = \mathbb{E}_{x_0, x_1, t}[\|v_\theta(x_t, t) - (x_1 - x_0)\|_2^2]$.

In the first stage, we generate a coarse voxel grid from the input condition by producing a low-resolution latent structure through a lightweight VAE trained to reconstruct voxel grids, following [65]. The pretrained VAE decoder efficiently upsamples this latent grid into a full-resolution voxel representation. In the second stage, we generate local SLAT features (z_i) for each active voxel (v_i) using the predicted voxel grid and conditioning input. Finally, the pretrained parametric surface decoder transforms $\{(v_i, z_i)_{i=1}^N\}$ into the final parametric surface

$\mathcal{S}$. For conditional embedding, we use modality-specific encoders: DINOv2 [46] for images and PointNet++ [48] for point clouds. We use pretrained weights for DINOv2 while training PointNet++ jointly with the Flow Transformer models.

Text-to-CAD Generation. While our framework supports text conditioning, training a direct text-to-3D model typically leads to slow convergence [52] and low prompt fidelity [53], due to the lack of explicit spatial and geometric cues in textual input. Following standard text-to-3D practices [30, 52, 53, 70], we adopt a two-stage approach: *text-to-image* followed by *image-to-CAD*. We fine-tune Stable Diffusion 3.5-2B [10] on the CADCap-1M dataset to align its output image distribution with that of the *image-to-CAD* model [53]. The resulting images are then used to condition our pretrained *image-to-CAD* model. This significantly enhances prompt fidelity with reduced training cost. Notably, pretrained *text-to-image* models [10, 25] often fail to preserve numerically constrained features such as hole or gear-tooth counts, underscoring the importance of fine-tuning for accurate *text-to-image* generation for the CAD domain.

5 Experiments

We evaluate DreamCAD on conditional CAD generation from text, image, and point inputs in Section 5.1. Section 5.2 analyzes the caption quality of the CADCap-1M dataset, and Section 5.3 presents ablation studies on key architectural and training design choices.

Datasets. We curate over 1M high-quality CAD models from 10 public datasets (Table 1) to train the DreamCAD VAE, converting all BRep-only datasets to meshes using OpenCascade. Although large-scale datasets like ABC and Automate contain BReps, training directly on BReps is impractical at scale as BRep topology is discrete and non-differentiable. We filter low-quality and trivial primitives from ABC [23] and Automate [19], splitting into 95% training, 2.5% validation, and 2.5% testing. For *text-to-CAD*, we use CADCap-1M and MARVEL [53] Level-5 captions for ShapeNet [4], Toys4K [54], and ABO [13]. For *image-to-CAD*, we render four orthographic views (front, back, left, right) in Blender [3] and randomly sample one view per epoch. For *point-to-CAD*, we normalize point clouds to $[-0.5, 0.5]$ and augment with surface normals in 50% of batches. We discuss more details in the supplementary.

Implementation. We train DreamCAD VAE on $8 \times$ NVIDIA H200 GPUs for 700k steps for 3 weeks with batch size 32 using mixed-precision, AdamW [36] with learning rate 5×10^{-5} and weight decay 1×10^{-4} . Both VAE encoder and decoder contain 8 Transformer layers with voxel latent dimension $z_i = 8$. We empirically set the loss weights in Eq. 5 to $\lambda_{cd} = 10^2$, $\lambda_{g1} = 5 \times 10^{-3}$, and $\lambda_{lp} = 1$. During training, points for CD loss increase from 16K to 100K via sigmoid scheduling and tessellation resolution r increases from (4, 4) to (16, 16). We train both coarse and fine-grained flow Transformer decoders for 500k steps on 4 NVIDIA H100 GPUs with 10% condition dropout. Pretrained image embeddings have a dimensionality of 1536. During inference, we use 50 steps and set the classifier free guidance scale to 7.5. For Stable Diffusion fine-tuning (text to image), we




Fig. 5: Qualitative comparison on **Point2CAD** (Top-Right), **Image2CAD** (Bottom-Left) and **Text2CAD** (Right) tasks. For each task, the first four examples are from the ABC dataset, while the last two from Objaverse dataset. **X** indicates invalid models.

apply LoRA [16] with rank and $\alpha=4$, trained for 300k steps on 4 H100 GPUs for 1 day. Inference takes ~ 15 s for *image-* and *point-to-CAD*, and ~ 30 s for *text-to-CAD* on a single H100 GPU.

5.1 Multimodal Generation Evaluation

Experimental Setup. We evaluate DreamCAD on two datasets: ABC [23] and Objaverse [7], each containing 15K samples. ABC serves as the in-distribution benchmark, while Objaverse is used for out-of-distribution (OOD) evaluation to measure generalization. Because Objaverse includes many free-form and organic objects uncharacteristic of CAD geometry, we filter its test set using MARVEL [53] captions containing CAD-specific keywords.

Baselines. As the first CAD generative framework trained on large-scale unstructured 3D data, DreamCAD has no direct counterparts. We therefore compare against representative design-history and UV-based methods across all tasks. For *text-to-CAD*, we include Text2CAD [21], Text2CQ [66], Cadrille [24], DeepCAD [62], CADFusion [57], NURBGen [55] and BRepDiff [27]. For CAD-Fusion [57], we generate 5 outputs per sample as per the official implementation. We observe this drastically reduces the invalidity ratio. For the *image-to-CAD* task, we compare DreamCAD against BRepDiff and Cadrille while for the *point-to-CAD* we additionally include DeepCAD and CAD-Recode [50]. DeepCAD is trained for 100 epochs on both *point* and *text-to-CAD* tasks on DeepCAD and Text2CAD datasets, while BRepDiff is trained for *image-to-cad* for 6k epochs on

Table 2: Quantitative comparison on Point2CAD, Text2CAD, and Img2CAD tasks over the ABC and Objaverse datasets. F1 and IR are scaled by 10^2 , while CD, JSD, and MMD are scaled by 10^3 . For *text* and *image-to-CAD*, GPT and User ratings measure visual alignment.

Task	Models	ABC						Objaverse							
		F1 ↑	NC ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓	F1 ↑	NC ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓
Point2CAD	DeepCAD [62]	19.31	0.49	51.10	0.37	783.94	29.63	11.01	7.05	0.48	320.33	0.41	855.14	34.62	10.88
	CAD-Recode [50]	75.99	0.79	3.73	0.13	271.89	2.94	15.39	53.24	0.66	7.92	0.19	479.50	6.27	15.53
	Cadrille [24]	78.86	0.80	2.98	0.12	236.10	2.51	5.84	57.49	0.67	6.28	0.17	445.23	5.24	10.81
	DreamCAD (Ours)	92.12	0.94	0.93	0.06	96.13	0.84	0.00	87.31	0.89	1.25	0.11	189.12	1.86	0.00
Img2CAD		GPT ↑	User ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓	GPT ↑	User ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓
	Cadrille [24]	7.75	5.34	111.50	0.49	909.92	68.97	14.19	1.10	0.45	99.91	0.48	913.24	58.02	27.16
	BRepDiff [27]	16.13	17.63	20.69	0.28	662.97	13.89	0.11	18.12	16.63	57.51	0.43	875.68	27.83	0.96
	DreamCAD (Ours)	76.12	77.03	4.12	0.17	412.31	6.31	0.00	80.78	82.92	20.16	0.27	541.81	13.41	0.006
Text2CAD		GPT ↑	User ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓	GPT ↑	User ↑	CD ↓	HD ↓	JSD ↓	MMD ↓	IR ↓
	DeepCAD [62]	0.49	0.40	86.54	0.44	887.69	37.66	39.02	0.00	0.00	80.68	0.45	903.40	37.34	41.9
	Text2CAD [21]	1.11	2.46	82.22	0.41	852.68	40.65	6.43	0.22	0.36	93.96	0.47	896.22	44.79	9.37
	Cadrille [24]	0.91	0.34	155.80	0.53	957.51	96.89	4.44	0.04	0.00	162.16	0.55	961.53	76.39	6.55
	Text2CQ [66] (Qwen3B)	0.01	0.00	68.15	0.39	829.72	37.50	12.52	0.00	0.00	83.43	0.44	890.89	47.62	14.4
	Text2CQ (GPT2L)	0.02	0.00	71.27	0.39	838.54	40.35	34.05	0.00	0.00	84.85	0.47	891.96	53.70	64.25
	Text2CQ (CodeGPT)	0.94	1.00	77.91	0.41	850.23	43.71	32.56	0.24	0.12	86.75	0.46	879.57	58.48	25.14
	CADfusion [57]	2.35	2.76	56.36	0.31	789.12	27.54	9.21	2.67	2.35	81.03	0.43	853.43	58.11	12.4
	NURBGen [55]	4.21	4.44	50.84	0.32	800.46	29.53	9.12	8.18	7.28	73.54	0.41	839.08	38.69	2.67
	BRepDiff [27]	4.34	3.20	54.12	0.38	812.31	34.72	1.12	6.33	6.41	74.32	0.38	808.19	41.13	2.31
	DreamCAD (Ours)	85.62	85.40	20.32	0.14	734.92	19.43	0.00	82.32	83.48	34.61	0.28	698.21	28.14	0.00

ABC. Since the training schemes introduce a data-scale disparity, we additionally report results on the DeepCAD dataset in the supplementary.

Metrics. We evaluate all tasks using both geometric and perceptual metrics. Geometric fidelity is assessed by Chamfer Distance (CD), Hausdorff Distance (HD), Jensen–Shannon Divergence (JSD), Minimum Matching Distance (MMD), Normal Consistency (NC), F1 score, and Invalidity Ratio (IR), which quantifies the fraction of non-convertible BRep outputs. All geometric metrics are computed on 8192 uniformly sampled points normalized within a unit cube centered at the origin. For *text*- and *image-to-CAD*, visual alignment is measured on 5k and 1k samples through GPT-5 [63] and user studies by 14 CAD-experts, respectively. In both settings, GPT-5 and human evaluators are shown multi-view renderings of reconstructions from all baselines and DreamCAD, and asked to select the model best matching the input (*text* or *image*). The same expert group is used consistently across all user studies.

Results. As shown in Table 2, DreamCAD achieves state-of-the-art results across all three modalities in *point-to-CAD*, *image-to-CAD*, and *text-to-CAD*. In the easier *point-to-CAD* task, it outperforms baselines on ABC by a large margin, reducing CD by up to 68% and 75% over Cadrille and CAD-Recode and improving F1 scores by 17% and 21%, respectively, while attaining a zero invalidity ratio. Similar gains are observed across other metrics and the Objaverse dataset as well. As illustrated in Figure 5.A, DreamCAD accurately reconstructs complex geometries such as gear wheels (Ex. 3) and chair (Ex. 5), plane (Ex. 6), whereas prior methods capture only the coarse shapes.

For *image-to-CAD*, DreamCAD achieves over 75% preference in both GPT and human evaluations on ABC and Objaverse, improving CD by 80% and 58% and

Table 3: Ablation studies on regularization (left), and voxel resolution choice (right).**(a)** Impact of regularizers on VAE performance.**(b)** Ablation on voxel-grid resolution

Model	CD ($\times 10^3$) ↓	Lap ($\times 10^3$) ↓	HD ↓
No Regularization	0.0210	0.0073	0.02
+ G1	0.0230	0.0064	0.022
+ Lap	0.0225	0.0022	0.022
+ G1 + Lap	0.0259	0.0020	0.024

Voxel-Grid	24	32	48	64
#patches	1434.32	2546.18	7720.48	10179.34
CD ($\times 10^3$)	0.0231	0.011	0.0109	0.0105

MMD by 54% and 52% over BRepDiff, respectively. As shown in Figure 5.B, DreamCAD produces high-fidelity reconstructions from single images via its coarse-to-fine pipeline and achieves the lowest invalidity ratio, failing only when no active voxels are generated in the coarse stage.

For the most challenging *text-to-CAD* task, DreamCAD attains over 80% preference in both GPT and user studies, substantially outperforming all baselines. The closest competitor, BRepDiff, which takes as input the same images generated by our finetuned SD 3.5, remains far behind in reconstruction quality. Geometrically, DreamCAD achieves a 62% reduction in CD with a zero invalidity ratio in the ABC dataset. As shown in Figure 5.C, it exhibits strong prompt fidelity, accurately reconstructing intricate shapes (Ex. 1: table, Ex. 6: camera) and numerically constrained features (Ex. 4: hole counts). Design-history-based models fail beyond simple primitives, while NURBGen’s lower prompt fidelity demonstrates that end-to-end text-to-CAD without visual grounding struggles to capture precise geometric details. Furthermore, BRepDiff’s post-processing stage occasionally produces disconnected BRep faces and visible spike artifacts from misplaced UV grid points, a limitation of grid-based representations that prevents reliable C^0 continuity.

5.2 Caption Quality

As **CADCap-1M** is the first large-scale captioning dataset for CAD models, no existing benchmark enables direct comparison. We therefore evaluate caption quality through both user studies and GPT-5 assessment on 1k and 5k samples respectively. Given four rendered views, metadata, and the caption, evaluators rate both geometric and semantic accuracy. Overall, 95.8% (user) and 98.31% (GPT-5) of captions are judged correct, including precise identification of part names and hole counts. This validates the reliability of our metadata-augmented prompting. We provide more statistical analyses in the supplementary material.

5.3 Ablation Study

Regularizations. We analyze the impact of **G1** and **Laplacian** regularizers on VAE reconstructions. We train the VAE from scratch on 300K samples and evaluating on 15K ABC test shapes (Table 3a and Figure 6). Without regularization, CD is minimized aggressively but produces rough surfaces with spike artifacts, reflected in the highest Laplacian value. Adding **G1** or **Laplacian** individually improves smoothness, with **Laplacian** better preserving curvature.




Fig. 6: Examples showing VAE reconstructions for different regularizers.




Fig. 7: Left: Captions from GPT-5 with and without metadata-augmented prompting. Right: Images from pretrained vs. fine-tuned SD-3.5 given the same prompts.

Combining both gives the lowest Laplacian loss (0.0020) and smoothest surfaces while maintaining strong geometric accuracy.

Voxel-Grid Resolution. We ablate the voxel-grid resolution used to initialize the parametric surfaces (Table 3.b). Since retraining the VAE at each resolution is expensive, we instead optimize 1000 randomly sampled training meshes. For each resolution, we voxelize each mesh, initialize the parametric surfaces, and optimize for 2000 epochs with learning rate 10^{-4} and loss from Eq 5, supervised by ground-truth points. Increasing resolution from 32 to 48 and 64 increases patch count by $3\times$ and $4\times$, while improving CD by only 1% and 5%. We therefore adopt resolution 32 as the best quality–efficiency trade-off.

Metadata Augmentation. We caption 1,000 random samples with and without metadata using GPT-5. Users preferred captions with metadata in 80.3% of cases, with more accurate part names and hole counts (Figure 7 Left).

Fine-Tuning Text-to-Image Model. We compare 500 images generated by pretrained and fine-tuned Stable Diffusion models. User studies show that the fine-tuned model is preferred in 75.6% of cases, demonstrating a notable improvement in prompt fidelity (Figure 7 -Right).

5.4 Application: CAD Topology Recovery

While DreamCAD’s outputs are editable via control points and weights, they still lack complete CAD topology. This remains an open challenge and forms




Fig. 8: (Left) Topology recovery from patch-based outputs to hybrid NURBS CAD representation [55]. (Right) For multimodal inputs, DreamCAD reconstructions and the corresponding recovered topology.

the basis of our future work. To explore feasibility, we fine-tune Qwen3-4B [68] with LoRA [16] on 50K samples to convert patch-based outputs into structured NURBS representations with semantic topology, following NURBGen [55]. Each patch is represented as 16 control points with weights, encoded via a Transformer encoder, and passed to Qwen3 for NURBS sequence prediction (Figure 8 - Left). Since patch counts remain low (e.g., $N_f = 2364$ in Figure 8), training is computationally tractable. Evaluated on 600 test samples across all three tasks (point-, image-, and text-to-CAD, 200 each), this generates 99.2% valid CAD models (Figure 8 - Right) with $CD = 0.17 \times 10^{-3}$. These results suggest that accurate and compact geometric reconstruction can provide a strong foundation for topology recovery for production-ready CAD generation. Full details are in the supplementary.

6 Conclusion

We present **DreamCAD**, a multi-modal generative framework that produces editable parametric surfaces directly from point-level supervision. Our goal is to address the long-standing challenge of generalizability in CAD generation. DreamCAD leverages a parametric patch-based CAD representation that supports differentiable mesh generation and enables point-based supervision for 3D shape synthesis. By removing the need for CAD-specific ground-truth annotations, our framework scales to large-scale 3D datasets for CAD geometry generation. In addition, we introduce **CADCap-1M**, the largest captioning dataset with high-quality textual descriptions generated using GPT-5 for advancing text-to-CAD research. Trained on 1M+ 3D meshes curated from 10 public datasets, DreamCAD demonstrates strong generalization across text, image, and point-conditioned generation tasks. While complete CAD topology recovery remains a hard and open problem, we view DreamCAD as establishing the geometric foundation upon which this challenging next stage can be built.

References

1. Ali, S.A., Khan, M.S., Stricker, D.: Brep boundary and junction detection for cad reverse engineering. In: ICML. IEEE (2024)
2. Autodesk Inc.: Autodesk Fusion 360: Integrated cad, cam, and cae software. <https://www.autodesk.com/products/fusion-360/> (2024)
3. Blender Online Community: Blender – a 3d modelling and rendering package. <https://www.blender.org/> (2025)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv:1512.03012 (2015)
5. Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Yago Vicente, T.F., Dideriksen, T., Arora, H., Guillaumin, M., Malik, J.: Abo: Dataset and benchmarks for real-world 3d object understanding. In: CVPR (2022)
6. Contero, M., Pérez-López, D., Company, P., Camba, J.D.: A quantitative analysis of parametric cad model complexity and its relationship to perceived modeling complexity. *Advanced Engineering Informatics* **56** (2023). <https://doi.org/https://doi.org/10.1016/j.aei.2023.101970>, <https://www.sciencedirect.com/science/article/pii/S1474034623000988>
7. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: CVPR (2023)
8. Du, W.H., Schmitt, F.J.: On the g1 continuity of piecewise bézier surfaces: a review with new results. *Computer-Aided Design* (1990)
9. Dupont, E., Cherenkova, K., Mallis, D., Gusev, G., Kacem, A., Aouada, D.: Transcad: A hierarchical transformer for cad sequence inference from point clouds. In: ECCV (2024)
10. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: ICML (2024)
11. Farin, G.: *Curves and surfaces for CAGD: a practical guide*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 5th edn. (2001)
12. Foley, J.D., van Dam, A., Feiner, S.K., Hughes, J.F.: *Computer graphics: principles and practice* (2nd ed.). Addison-Wesley Longman Publishing Co., Inc., USA (1990)
13. Fu, H., Jia, R., Gao, L., Gong, M., Zhao, B., Maybank, S., Tao, D.: 3d-future: 3d furniture shape with texture. *IJCV* (2021)
14. Govindarajan, P., Baldelli, D., Pathak, J., Fournier, Q., Chandar, S.: Cadmium: Fine-tuning code language models for text-driven sequential cad design. arXiv:2507.09792 (2025)
15. Guo, H., Liu, S., Pan, H., Liu, Y., Tong, X., Guo, B.: Complexgen: Cad reconstruction by b-rep chain complex generation. *SIGGRAPH* (2022)
16. Hu, E.J., yelong shen, Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: LoRA: Low-rank adaptation of large language models. In: *ICLR* (2022)
17. Hunde, B., Woldeyohannes, A.: Future prospects of computer-aided design (cad) – a review from the perspective of artificial intelligence (ai), extended reality, and 3d printing. *Results in Engineering* (2022)
18. Jayaraman, P.K., Sanghi, A., Lambourne, J.G., Willis, K.D., Davies, T., Shayani, H., Morris, N.: Uv-net: Learning from boundary representations. In: *CVPR* (2021)
19. Jones, B., Hildreth, D., Chen, D., Baran, I., Kim, V.G., Schulz, A.: Automate: A dataset and learning approach for automatic mating of cad assemblies. *SIGGRAPH* (2021)

20. Khan, M.S., Dupont, E., Ali, S.A., Cherenkova, K., Kacem, A., Aouada, D.: Cad-signet: Cad language inference from point clouds using layer-wise sketch instance guided attention. In: CVPR (2024)
21. Khan, M.S., Sinha, S., Uddin, S.T., Stricker, D., Ali, S.A., Afzal, M.Z.: Text2cad: Generating sequential CAD designs from beginner-to-expert level text prompts. In: NeurIPS (2024)
22. Khanna, M., Mao, Y., Jiang, H., Haresh, S., Shacklett, B., Batra, D., Clegg, A., Undersander, E., Chang, A.X., Savva, M.: Habitat synthetic scenes dataset (hssd-200): An analysis of 3d scene scale and realism tradeoffs for objectgoal navigation. In: CVPR (2024)
23. Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M., Zorin, D., Panozzo, D.: Abc: A big cad model dataset for geometric deep learning. In: CVPR (2019)
24. Kolodiaznyi, M., Tarasov, D., Zhemchuzhnikov, D., Nikulin, A., Zisman, I., Vorontsova, A., Konushin, A., Kurenkov, V., Rukhovich, D.: cadrille: Multi-modal cad reconstruction with online reinforcement learning. arXiv:2505.22914 (2025)
25. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux> (2024)
26. Lambourne, J.G., Willis, K.D., Jayaraman, P.K., Sanghi, A., Meltzer, P., Shayani, H.: Brepnet: A topological message passing system for solid models. In: CVPR (June 2021)
27. Lee, M., Zhang, D., Jambon, C., Kim, Y.M.: Brepdiff: Single-stage b-rep diffusion model. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (2025)
28. Li, J.: Applications of computer-aided design technology in research, engineering and industry. Computer Aided Design: Technology, Types and Practical Applications (2012)
29. Li, J., Ma, W., Li, X., Lou, Y., Zhou, G., Zhou, X.: Cad-llama: Leveraging large language models for computer-aided design parametric 3d model generation. In: CVPR (2025)
30. Li, J., Tan, H., Zhang, K., Xu, Z., Luan, F., Xu, Y., Hong, Y., Sunkavalli, K., Shakhnarovich, G., Bi, S.: Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. In: ICLR (2024)
31. Li, J., Fu, Y., Chen, F.: Dtgbreppen: A novel b-rep generative model through decoupling topology and geometry. In: CVPR (2025)
32. Li, P., Guo, J., Zhang, X., Yan, D.M.: Secad-net: Self-supervised cad reconstruction by learning sketch-extrude operations. In: CVPR (2023)
33. Li, W., Zhang, X., Sun, Z., Qi, D., Li, H., Cheng, W., Cai, W., Wu, S., Liu, J., Wang, Z., et al.: Step1x-3d: Towards high-fidelity and controllable generation of textured 3d assets. arXiv:2505.07747 (2025)
34. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: ICLR (2023)
35. Liu, Y., Obukhov, A., Wegner, J.D., Schindler, K.: Point2cad: Reverse engineering cad models from 3d point clouds. In: CVPR (2024)
36. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019)
37. Luo, T., Rockwell, C., Lee, H., Johnson, J.: Scalable 3d captioning with pretrained models. In: NeurIPS (2023)
38. Lv, C., Bao, J.: Cadinstruct: A multimodal dataset for natural language-guided cad program synthesis. Computer-Aided Design (2025)
39. Mallis, D., Aziz, A.S., Dupont, E., Cherenkova, K., Karadeniz, A.S., Khan, M.S., Kacem, A., Gusev, G., Aouada, D.: Sharp challenge 2023: Solving cad history and

- parameters recovery from point clouds and 3d scans. overview, datasets, metrics, and baselines. In: ICCVw (2023)
40. Mildenhall, B., Srinivasan, P., Tancik, M., Barron, J., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
41. Milnor, J.W., Weaver, D.W.: Topology from the differentiable viewpoint, vol. 21. Princeton university press (1997)
42. Nealen, A., Igarashi, T., Sorkine, O., Alexa, M.: Laplacian mesh optimization. In: Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia (2006)
43. Onshape Inc.: Onshape: Cloud-based cad platform. <https://www.onshape.com/> (2024)
44. Open Cascade SAS: Opencascade technology (occt). <https://www.opencascade.com/open-cascade-technology/> (2024)
45. OpenAI: Gpt-5. <https://openai.com/gpt-5/> (2025)
46. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision. arXiv:2304.07193 (2023)
47. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
48. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. NeurIPS (2017)
49. Ren, D., Zheng, J., Cai, J., Li, J., Zhang, J.: Extrudenet: Unsupervised inverse sketch-and-extrude for shape parsing. In: ECCV (2022)
50. Rukhovich, D., Dupont, E., Mallis, D., Cherenkova, K., Kacem, A., Aouada, D.: Cad-recode: Reverse engineering cad code from point clouds. In: ICCV (2025)
51. Sharma, G., Liu, D., Maji, S., Kalogerakis, E., Chaudhuri, S., M  ch, R.: Parsenet: A parametric surface fitting network for 3d point clouds. In: ECCV (2020)
52. Siddiqui, Y., Monnier, T., Kokkinos, F., Kariya, M., Kleiman, Y., Garreau, E., Gafni, O., Neverova, N., Vedaldi, A., Shapovalov, R., Novotny, D.: Meta 3d assetgen: Text-to-mesh generation with high-quality geometry, texture, and pbr materials. In: NeurIPS (2024)
53. Sinha, S., Khan, M.S., Usama, M., Sam, S., Stricker, D., Ali, S.A., Afzal, M.Z.: Marvel-40m+: Multi-level visual elaboration for high-fidelity text-to-3d content creation. In: CVPR (2025)
54. Stojanov, S., Thai, A., Rehg, J.M.: Using shape to categorize: Low-shot learning with an explicit shape bias. In: CVPR (2021)
55. Usama, M., Khan, M.S., Stricker, D., Afzal, M.Z.: Nurbgen: High-fidelity text-to-cad generation through llm-driven nurbs modeling. In: AAAI (2026)
56. Uy, M.A., yu Chang, Y., Sung, M., Goel, P., Lambourne, J., Birdal, T., Guibas, L.: Point2cyl: Reverse engineering 3d objects from point clouds to extrusion cylinders. In: CVPR (2022)
57. Wang, R., Yuan, Y., Sun, S., Bian, J.: Text-to-cad generation through infusing visual feedback in large language models. In: ICML (2025)
58. Wang, S., Chen, C., Le, X., Xu, Q., Xu, L., Zhang, Y., Yang, J.: Cad-gpt: Synthesising cad construction sequence with spatial reasoning-enhanced multimodal llms. In: AAAI (2025)
59. Wang, Z.: 3d representation methods: A survey. arXiv:2410.06475 (2024)

60. Willis, K.D.D., Pu, Y., Luo, J., Chu, H., Du, T., Lambourne, J.G., Solar-Lezama, A., Matusik, W.: Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. SIGGRAPH (2021)
61. Worchel, M., Alexa, M.: Differentiable rendering of parametric geometry. SIGGRAPH (2023)
62. Wu, R., Xiao, C., Zheng, C.: Deepcad: A deep generative network for computer-aided design models. In: ICCV (2021)
63. Wu, T., Yang, G., Li, Z., Zhang, K., Liu, Z., Guibas, L., Lin, D., Wetzstein, G.: Gpt-4v(ision) is a human-aligned evaluator for text-to-3d generation. In: CVPR (2024)
64. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR (2015)
65. Xiang, J., Lv, Z., Xu, S., Deng, Y., Wang, R., Zhang, B., Chen, D., Tong, X., Yang, J.: Structured 3d latents for scalable and versatile 3d generation. In: CVPR (2025)
66. Xie, H., Ju, F.: Text-to-cadquery: A new paradigm for cad generation with scalable large model capabilities. arXiv:2505.06507 (2025)
67. Xu, X., Lambourne, J., Jayaraman, P., Wang, Z., Willis, K., Furukawa, Y.: Brep-gen: A b-rep generative diffusion model with structured latent geometry. SIGGRAPH (2024)
68. Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al.: Qwen3 technical report. In: Arxiv (2025)
69. You, Y., Uy, M.A., Han, J., Thomas, R., Zhang, H., Du, Y., Chen, H., Engelmann, F., You, S., Guibas, L.: Img2cad: Reverse engineering 3d cad models from images through vlm-assisted conditional factorization. arXiv:2408.01437 (2024)
70. Zhang, L., Wang, Z., Zhang, Q., Qiu, Q., Pang, A., Jiang, H., Yang, W., Xu, L., Yu, J.: Clay: A controllable large-scale generative model for creating high-quality 3d assets. SIGGRAPH (2024)
71. Zhao, Z., Lai, Z., Lin, Q., Zhao, Y., Liu, H., Yang, S., Feng, Y., Yang, M., Zhang, S., Yang, X., et al.: Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. CoRR (2025)
72. Zhou, S., Tang, T., Zhou, B.: Cadparser: a learning approach of sequence modeling for b-rep cad. In: IJCAI (2023)
73. Zou, Q., Wu, Y., Liu, Z., Xu, W., Gao, S.: Intelligent cad 2.0. Visual Informatics (2024)