

ReCAD: Reinforcement Learning Enhanced Parametric CAD Model Generation with Vision-Language Models

Jiahao Li, Yusheng Luo, Yunzhong Lou, Xiangdong Zhou*

College of Computer Science and Artificial Intelligence, Fudan University, Shanghai, China
Shanghai Key Laboratory of Data Science
{lijiahao25, ycluo24}@m.fudan.edu.cn, {yzlou20, xdzhou}@fudan.edu.cn

Abstract

We present ReCAD, a reinforcement learning (RL) framework that bootstraps pretrained large models (PLMs) to generate precise parametric computer-aided design (CAD) models from multimodal inputs by leveraging their inherent generative capabilities. With just access to simple functional interfaces (e.g., point coordinates), our approach enables the emergence of complex CAD operations (e.g., *pattern replication* and *mirror*). This stands in contrast to previous methods, which typically rely on knowledge injected through supervised fine-tuning (SFT), offer limited support for editability, and fail to exploit the strong generative priors of PLMs. Specifically, the ReCAD framework begins by fine-tuning vision-language models (VLMs) to equip them with basic CAD model generation capabilities, where we rewrite CAD scripts into parameterized code that is leveraged to generate accurate textual descriptions for supervision. Then, we propose a novel RL strategy that incorporates parameterized code as guidance to enhance the model’s reasoning on challenging questions. Furthermore, we employ a hierarchical primitive learning process to progressively teach structured and compositional skills under a unified reward function that ensures both geometric accuracy and semantic fidelity. ReCAD sets a new state-of-the-art in both text-to-CAD and image-to-CAD tasks, significantly improving geometric accuracy across in-distribution and out-of-distribution settings. In the image-to-CAD task, for instance, it reduces the mean Chamfer Distance from 73.47 to 29.61 (in-distribution) and from 272.06 to 80.23 (out-of-distribution), outperforming existing baselines by a substantial margin.

Introduction

Prototyping complex computer-aided design (CAD) models is a time-consuming process, as it often involves numerous parts and demands high precision, requiring designers to carefully craft each detail (Cherng et al. 1998; Robertson and Allen 2002). As a result, generative CAD modeling has attracted increasing attention from both the research and industrial communities (Daareyni et al. 2025). Traditional CAD generation models often employ encoder-decoder architectures (Wu, Xiao, and Zheng 2021; Xu et al. 2023), using modality-specific encoders (e.g., for point clouds (Ma

et al. 2024) or text (Khan et al. 2024b; Li et al. 2024)) and autoregressive decoders to produce CAD models. Recently, CAD generation has shifted towards using pretrained large models (PLMs) to create CAD models, which are represented in the form of code, JSON, or command sequences (Wang et al. 2025b; Li et al. 2025a; Xu et al. 2024).

Despite these advancements, the generation of high-precision CAD models from descriptions that encode dimensional and quantitative constraints remains a significant challenge. Ensuring that the generated models align with the specified original design intent is essential for facilitating subsequent reuse and editability (Martin 2023). Nevertheless, previous methodologies lack the capability to comprehend design intent, as they directly produce CAD sequences consisting solely of low-level parameters (e.g., point coordinates) with modeling commands (Li et al. 2025a; Wang et al. 2025a; Zhang et al. 2024). The adjustment of these parameters can readily lead to the creation of invalid geometries, such as non-closed loops (Chen et al. 2025).

At present, the majority of methodologies reconceptualize computer-aided design (CAD) modeling as an endeavor of code generation (Li et al. 2025a; Guan et al. 2025; Doris et al. 2025). This shift is primarily attributed to the potent capabilities of PLMs (Gao et al. 2023; Wong et al. 2023). Nevertheless, PLMs typically act as semantic interpreters (Yuan et al. 2025) in these works, as their ability to generate and understand CAD models largely relies on the knowledge gained during supervised fine-tuning (SFT). This limits the use of PLMs’ generative capabilities, leading to dependence on external knowledge and weak generalization.

Parametric CAD modeling inherently requires precise mathematical reasoning, symbolic manipulation, and the satisfaction of logical constraints (Zhao et al. 2025), which are essential for capturing design intent. Recent advances in reinforcement learning with verifiable rewards (RLVR) (Lambert et al. 2024) have demonstrated promising improvements in the reasoning abilities of PLMs, particularly in mathematics and code generation (Luo et al. 2025; Shao et al. 2024). To this end, we propose ReCAD, a reinforcement learning framework designed to enhance both the reasoning capability and generative performance of PLMs in the context of parametric CAD modeling.

Specifically, we first fine-tune vision-language models (VLMs) by converting CAD sequences into parameterized

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Visualization of the 3D model generated by our method (ReCAD-VL) based on a single image or textual description (abstract or precise). Our method enables the generation of complex and diverse parametric CAD models in code representation, which allows precise parameter control and facilitates easy editing and reuse.

code, which is then used to produce abstract and precise descriptions. These descriptions, together with the corresponding CAD images, are used to perform SFT on both text-to-CAD and image-to-CAD tasks. Building on this, we introduce a reinforcement learning stage using group relative policy optimization (GRPO) (Shao et al. 2024), leveraging parameterized code as off-policy guidance to enhance the model’s reasoning on inherently challenging questions. Additionally, a hierarchical primitive learning process is proposed to progressively teach compositional skills across abstraction levels. In order to achieve accurate generation, we develop a reward function that assesses both geometric precision and semantic fidelity, thereby ensuring reliable CAD generation.

We show that even with simple function interfaces that expose only the coordinate values of individual curves, while all other components serve merely as hierarchical wrappers, ReCAD enables the emergence of complex CAD operations such as *circular pattern* and *mirror*, which were previously unseen in earlier studies, as illustrated in Figure 1 and Figure 3. We also demonstrate that these emergent capabilities substantially enhance the model’s ability to generalize, leading to significantly improved geometric accuracy over the baseline methods under out-of-distribution (OOD) settings.

To summarize, our key contributions are as follows: (1) We introduce ReCAD, a novel reinforcement learning (RL) framework that enables the generation of precise and editable parametric CAD models. (2) We combine supervised fine-tuning with a novel RL strategy in which the parameterized CAD code itself acts as both an off-policy signal and complementary knowledge, boosting reasoning on complex questions. (3) We introduce a hierarchical primitive learning process that gradually builds the model’s ability to compose structured designs across abstraction levels under a unified reward for geometric and semantic fidelity. (4) Experimental results demonstrate that ReCAD outperforms existing methods in both text-to-CAD and image-to-CAD, while exhibiting remarkable generalization capabilities and zero-shot performance.

Related Work

CAD Generation and Reconstruction

Parametric CAD Sequence Modeling. The generation and reconstruction of CAD models have been a long-standing

research topic. Along with the emergence of large-scale datasets (Wu, Xiao, and Zheng 2021; Willis et al. 2021), prior works leverage representation learning methods (Wu, Xiao, and Zheng 2021; Xu et al. 2022, 2023), where CAD models are first encoded into latent representations, from which the corresponding CAD sequences are decoded. Some methods adopt transformer architectures (Khan et al. 2024a,b; Li et al. 2024), encoding inputs such as point clouds, text, or partial CAD models using modality-specific encoders, and then auto-regressively inferring the target CAD sequence. In addition, diffusion-based models (Ma et al. 2024; Li et al. 2025b,c; Chen et al. 2025; Chereddy and Femiani 2025) have been proposed to generate CAD sequences or directly synthesize B-Rep representations from inputs such as images, point clouds, or partial CAD models. However, these approaches often exhibit limited generalizability and struggle to generate precise CAD models.

PLM for CAD Modeling. PLMs have been adapted to CAD-related tasks. To leverage the prior knowledge in PLMs, many works elaborate task-specific datasets and subsequently apply supervised fine-tuning (SFT) to align PLMs with downstream applications involving CAD editing (Yuan et al. 2025; Zhang et al. 2024, 2025b) and generation (Wang et al. 2025b; Xu et al. 2024; Wang et al. 2025a; Wu et al. 2024; Rukhovich et al. 2025), where the models are conditioned on various input modalities such as images, text, or point clouds to produce CAD sequences. However, the ability of PLMs to understand and generate CAD sequences remains tightly coupled with the knowledge introduced during the SFT stage, limiting generative priors and hindering faithful modeling of design intent, such as parameter control. Another line of research explores constrained sketch generation to align outputs with design intent. Vitruvion (Seff et al. 2021) and SketchDNN (Chereddy and Femiani 2025) generate 2D sketches using token-based and diffusion approaches, respectively, while another work (Casey et al. 2025) adapts alignment techniques from LLMs. However, these methods are limited to 2D sketches, struggle with complex designs, and rely on external constraint solvers.

Reinforcement Learning for PLM. The paradigm of reinforcement learning with verifiable rewards (RLVR) has emerged as a pivotal technique for augmenting the reasoning capabilities of Pre-trained Language Models (PLMs). Unlike earlier methods based on preferences (Ouyang et al.

2022; Rafailov et al. 2023), RLVR leverages verifiable signals, yielding strong results on math, code (Lambert et al. 2024; Shao et al. 2024), and multimodal reasoning benchmarks (Shen et al. 2025; Zhan et al. 2025). On-policy methods like PPO (Schulman et al. 2017) and GRPO (Shao et al. 2024) ensure stable training but hinder exploration. LUFFY (Yan et al. 2025) integrates expert off-policy reasoning into advantage estimation, while another work (Zhang et al. 2025a) rewrites expert trajectories using the policy model to reduce distribution mismatch. However, obtaining such high-quality Chain-of-Thought (CoT) annotations is both resource-intensive and time-consuming. Related to our work, CAD-Coder (Guan et al. 2025) performs SFT with CoT annotations and RL using expert-level descriptions from the Text2CAD (Khan et al. 2024b) dataset, which contains overparameterized low-level details such as precise coordinates. While this improves CAD generation accuracy, it limits the model’s ability to infer low-level attributes from high-level design intent.

Method

We present ReCAD, a framework that leverages VLMs and RL to generate precise parametric CAD models from a single image or text. We first introduce the problem and notation, then describe a two-stage training process: supervised fine-tuning (SFT) and reinforcement learning (RL).

Problem Definition and Notation

ReCAD takes as input a CAD image I or a textual description T , either detailed (T^D) or abstract (T^A), and translates it into parametric CAD code \hat{C} representing the final model.

Sketch-Extrude Paradigm. We build upon the widely used *sketch-extrude* (SE) design paradigm, where a CAD model D is composed of multiple SE pairs, denoted as MSE . Each sketch S contains multiple faces; each face F consists of one or more loops L ; and each loop is a closed path formed by one or more curves. Based on this hierarchical structure, we define five levels of hierarchical primitives:

$$\mathcal{P} = \{L, F, S, SE, MSE\}, \quad (1)$$

where $P \in \mathcal{P}$ denotes any primitive in the hierarchy. A complete CAD model can be represented by a primitive P when $P = SE$ or $P = MSE$.

CAD Sequence Representation. Instead of adopting original CAD sequence representations (e.g., Line 128 128) or translating CAD sequences into CADQuery (CadQuery 2025) code, we introduce a lightweight set of function interfaces tailored to \mathcal{P} . At the lowest level, curves are defined via coordinate-based interfaces, and higher-level primitives P are progressively wrapped to ultimately form an MSE object that represents a complete CAD model. We define a function $f(\cdot)$ that converts a primitive P into our interface-based CAD code C , as follows:

$$C = f(P). \quad (2)$$

Supervised Fine-Tuning on Parameterized Code

Although reinforcement learning with verifiable rewards (RLVR) has shown promising results, it remains insufficient

for encouraging capabilities beyond the inherent limitations of the base model. Due to the limited proficiency of existing VLMs in CAD-specific tasks (Huang et al. 2025), we first employ a supervised fine-tuning (SFT) stage to establish foundational capabilities in parametric CAD code generation.

Parameterized Code. Prior work (Li et al. 2025a) has demonstrated that converting CAD sequences into code representations enhances model performance in CAD generation. However, such naive conversions fail to support modular coding practices and often produce inflexible, overfit-prone representations (referred to as hardcoded CAD code), which may limit their ability to generalize to out-of-distribution (OOD) samples.

To this end, we rewrite hardcoded CAD code into parameterized code using a VLM, aiming to unleash the LLM’s strong prior in programmatic code generation and to provide a more flexible representation of CAD code. For a primitive P , we first obtain its hardcoded code $C = f(P)$ and render a canonical-view image I of P . Specifically, we render a canonical view if $P \in \{SE, MSE\}$, and a top-down view otherwise. The image-code pair (I, C) is then fed into the VLM to sample N candidate parameterized codes:

$$\{\hat{C}^i\}_{i=1}^N = \text{VLM}(I, C). \quad (3)$$

To select high-quality candidates, we render each \hat{C}^i into an image \hat{I}_i and use a pretrained DINOv2 (Oquab et al. 2023) model as an encoder $E(\cdot)$ to compute the cosine similarity between $E(\hat{I}_i)$ and $E(I)$. We retain only candidates who satisfy:

$$\mathcal{C} = \left\{ \hat{C}^i \mid \cos(E(\hat{I}_i), E(I)) > \tau_s \right\}, \quad (4)$$

where τ_s is a similarity threshold. If no valid candidates exist, we revert to the original code: $\mathcal{C} = \{C\}$

Text Description Generation. To generate textual descriptions for CAD models, existing methods often rely on VLMs and complex annotation pipelines. However, these methods typically struggle to capture precise information (e.g., scale, quantity), and the resulting descriptions, even when incorporating CAD sequence information, tend to be overly verbose with redundant parameter details (Khan et al. 2024b); as such, low-level information is difficult for VLMs to interpret.

Instead, we leverage the parameterized code \mathcal{C} , which inherently contains semantically meaningful and precise information about both scale and quantity. Such information is explicitly encoded and is thus easily interpretable by VLMs. By feeding both the parameterized code and image I into a VLM, we obtain textual descriptions:

$$T^A, T^D = \text{VLM}(I, \hat{C}^i), \quad (5)$$

where T^A and T^D denote abstract and detailed descriptions, respectively, and $\hat{C}^i \in \mathcal{C}$.

Supervised Fine-Tuning. We employ the standard causal language modeling (CLM) objective to continuously fine-tune the LLM on two tasks: text-to-CAD and image-to-CAD. Formally, given the dataset $\mathcal{D}_{\text{SFT}} = \{q_i, \mathcal{C}_i\}_{i=1}^N$,

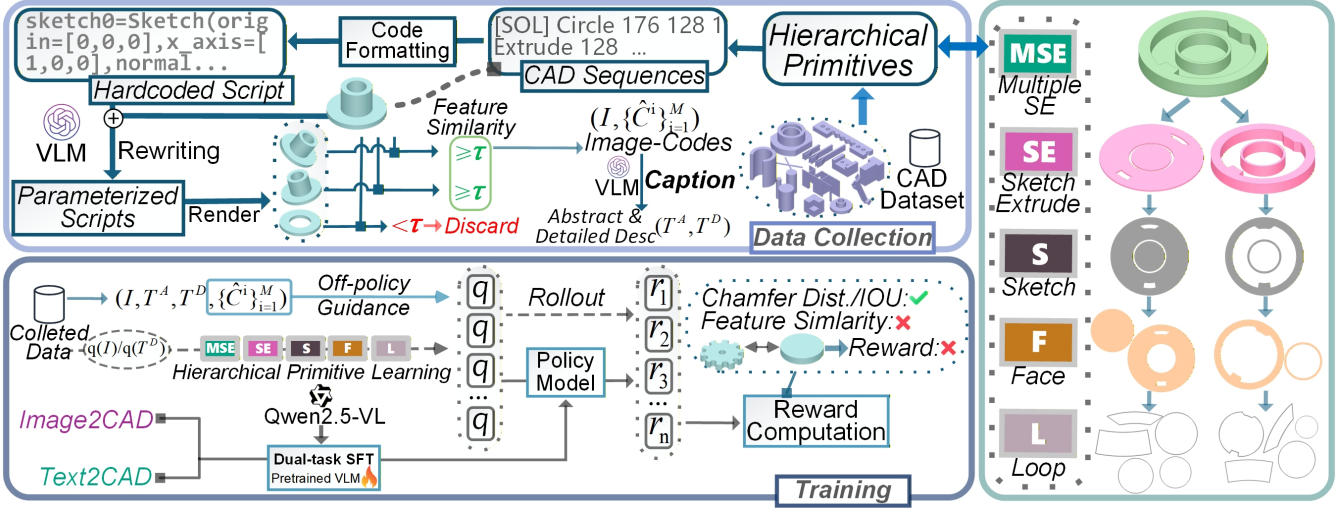


Figure 2: Overview of the proposed ReCAD framework. (1) A CAD model is represented as a primitive hierarchy, where each primitive is converted into parameterized code, which is further leveraged to generate precise textual descriptions. (2) We first perform supervised fine-tuning (SFT) for basic CAD code generation, then optimize the model via reinforcement learning guided by parameterized code, with reward functions enforcing both geometric fidelity and feature-level consistency.

where $q_i = T_i^D/T_i^A$ is for the text-to-CAD task and $q_i = I_i$ is for the image-to-CAD task, we optimize the LLM via maximum likelihood:

$$\mathcal{L}(\mathcal{D}_{\text{SFT}}) = - \sum_{i=1}^N \log \pi_{\theta}(\mathcal{Y}_i \in \mathcal{C}_i \mid \mathcal{X}_i = q_i). \quad (6)$$

Here, π_{θ} denotes the LLM with trainable parameters θ , and $\mathcal{X}_i, \mathcal{Y}_i$ represent the input and output, respectively.

To retain the LLM’s general abilities, we augment the training corpus with a small proportion of data sampled from *UltraChat* and *OpenCodeReasoning*. After fine-tuning, the model acquires basic CAD code generation abilities, and we name this model ReCAD-Base.

Reinforcement Learning Under Guidance

We employ RL after SFT to enhance the model’s generalization ability and its capability to generate precise geometry. Similar to the SFT stage, we perform RL on both text-to-CAD and image-to-CAD tasks and refer to the trained model as ReCAD-VL. In contrast to the SFT stage, the text-to-CAD input in the RL stage includes only the precise description T^D , since one abstract description T^A may map to multiple valid geometries, complicating reward design. We employ group relative policy optimization (GRPO), with ReCAD-Base serving as the base model. The GRPO objective is defined as:

$$\mathcal{J}(\theta) = \frac{1}{N} \sum_{i=1}^N \frac{1}{|\tau_i|} \sum_{t=1}^{|\tau_i|} \text{CLIP}(r_{i,t}(\theta, q), A_i, \epsilon) - \beta \mathbb{D}_{\text{KL}}[\pi_{\theta} \parallel \pi_{\text{ref}}], \quad (7)$$

where each question q is associated with a set of sampled solutions $\{\tau_i\}_{i=1}^N$ generated from $\pi_{\theta_{\text{old}}}$. The advantage A_i is computed using normalized rewards within the group:

$$A_i = \frac{R(\tau_i) - \text{mean}(\{R(\tau_j)\}_{j=1}^N)}{\text{std}(\{R(\tau_j)\}_{j=1}^N)}, \quad (8)$$

where $R(\cdot)$ is the reward function. The clipped surrogate objective $\text{CLIP}(r_{i,t}(\theta, q), A_i, \epsilon)$ is defined as:

$$\text{CLIP}(r, A, \epsilon) = \min[r \cdot A, \text{clip}(r, 1 - \epsilon, 1 + \epsilon) \cdot A], \quad (9)$$

where $r_{i,t}(\theta, q) = \pi_{\theta}(\tau_{i,t} \mid q, \tau_{i,<t}) / \pi_{\theta_{\text{old}}}(\tau_{i,t} \mid q, \tau_{i,<t})$ denotes the importance sampling ratio.

Learn Under Guidance. Despite the success of RLVR, its on-policy nature restricts learning to self-generated outputs and fundamentally limits it to the base LLM. Inspired by the strength of LLMs in few-shot in-context learning (Dong et al. 2022), we treat the rewritten parameterized codes \mathcal{C} as off-policy guidance during the rollout, providing complementary knowledge that helps the model develop coherent reasoning guided by correct code. The objective of this process is defined as follows:

$$\begin{aligned} \hat{\mathcal{J}}(\theta; \mathcal{C}) &= \frac{1}{N - |\mathcal{C}|} \sum_{i=1}^{N - |\mathcal{C}|} \frac{1}{|\tau_i|} \sum_{t=1}^{|\tau_i|} \text{CLIP}(r_{i,t}(\theta, q), A_i, \epsilon) \\ &+ \frac{1}{|\mathcal{C}|} \sum_{j=1}^{|\mathcal{C}|} \frac{1}{|\tau_j|} \sum_{t=1}^{|\tau_j|} \text{CLIP}(\hat{r}_{j,t}(\theta, q, C_j), A_j, \epsilon) \\ &- \beta \mathbb{D}_{\text{KL}}[\pi_{\theta} \parallel \pi_{\text{ref}}], \end{aligned} \quad (10)$$

where $C_j \in \mathcal{C}$ denotes the j -th parameterized code used for guidance, and τ_j is the corresponding solution generated by $\pi_{\theta_{\text{old}}}$ conditioned on the question q and C_j . The importance sampling ratio is defined as $\hat{r}_{j,t}(\theta, q, C_j) = \pi_{\theta}(\tau_{j,t} \mid q, \tau_{j,<t}) / \pi_{\theta_{\text{old}}}(\tau_{j,t} \mid q, C_j, \tau_{j,<t})$. The advantage values A are computed as in Equ. (8) over the group of solutions generated with and without guidance, i.e., τ_j and τ_i , respectively.

We define the RL training dataset as $\mathcal{D}_{\text{RL}} = \{q_i, \mathcal{C}_i\}_{i=1}^N$. Before RL training, we identify difficult questions by prompting the model with each q_i , sampling N solutions, and computing the maximum reward $\max\{R(q_i)\}$. If the

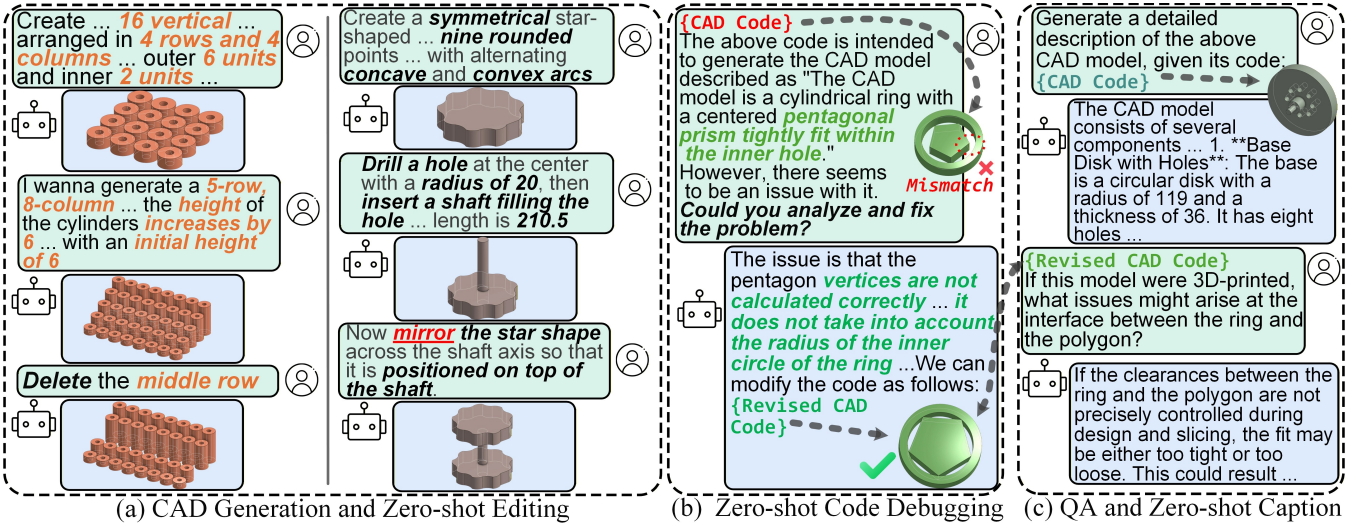


Figure 3: Dialog results on CAD generation and zero-shot CAD-related tasks. Despite being trained exclusively on CAD generation tasks, ReCAD-VL exhibits impressive zero-shot performance across multiple CAD-related tasks, including understanding, editing, debugging, and descriptive captioning.

reward falls below a predefined threshold τ_h , we identify q_i as a hard question and use the guidance-based objective $\hat{\mathcal{J}}(\theta; C_i)$. Otherwise, we apply the standard GRPO objective $\mathcal{J}(\theta)$. Thus, the final training objective is:

$$\mathcal{L}_{\text{RL}}(\theta) = \mathbb{E}_{(q_i, C_i) \sim \mathcal{D}_{\text{RL}}} \left[\mathbf{1}_{\text{hard}}(q_i) \cdot \hat{\mathcal{J}}(\theta; C_i) + (1 - \mathbf{1}_{\text{hard}}(q_i)) \cdot \mathcal{J}(\theta) \right], \quad (11)$$

where $\mathbf{1}_{\text{hard}}(q_i)$ is an indicator function that evaluates to 1 if the maximum reward for question q_i falls below a predefined threshold τ_h , and 0 otherwise:

$$\mathbf{1}_{\text{hard}}(q_i) = \begin{cases} 1 & \text{if } \max\{R(q_i)\} < \tau_h \\ 0 & \text{otherwise} \end{cases}. \quad (12)$$

Hierarchical Primitive Learning. We introduce hierarchical primitive learning (HPL) as a curriculum learning strategy tailored to the CAD generation task. Curriculum learning (Bengio et al. 2009; Deng et al. 2025) is a widely adopted strategy in complex task domains, wherein the learning process is organized into a sequence of increasingly difficult examples. This paradigm mirrors the human learning process (Elman 1993), enabling models to progressively acquire fundamental skills before tackling more intricate challenges. In our setting, CAD models follow the sketch-extrude (SE) paradigm, which naturally forms a hierarchical structure from curves to loops, faces, sketches, and extrusions. Based on the defined hierarchical primitives \mathcal{P} (Equ. (1)), we design a curriculum that follows the inherent structure of CAD models, progressing through L , F , S , SE , and MSE . Each stage introduces increasing complexity and builds upon the structure of the preceding one, allowing the model to acquire compositional skills in a gradual and structured manner. For each level, we further organize training samples by difficulty, which is defined by the number of

curves involved. The model learns in a simple-to-complex order, starting with primitives composed of fewer curves and gradually moving to those with richer geometric structures.

Reward Design. To ensure that the generated CAD models are both geometrically accurate and visually faithful, we design a reward function that combines intersection-over-union under optimal alignment (Doris et al. 2025), denoted as IOU_{best} , with feature-level similarity:

$$R(y_\pi, \Omega) = \lambda_1 \cdot \min\{\text{IOU}_{\text{best}}(\hat{\Omega}, \Omega), \phi(\text{sim}(\hat{I}, I), \tau)\} + \lambda_2 \cdot R_f(y_\pi), \quad (13)$$

where y_π denotes the model-generated output sequence. The geometric consistency between the predicted geometry $\hat{\Omega}$ and ground-truth Ω is measured by $\text{IOU}_{\text{best}}(\hat{\Omega}, \Omega)$, which computes the IOU and selects the alignment with maximal overlap. The image similarity $\text{sim}(\hat{I}, I) = \cos(E(\hat{I}), E(I))$ is the cosine similarity between the rendered images of $\hat{\Omega}$ and Ω . The function $\phi(s, \tau) = \max(0, (s - \tau)/(1 - \tau))$ applies thresholded linear scaling, with τ set to 0.55 in our experiments. Scalars λ_1 and λ_2 are weighting coefficients. The formatting reward $R_f(y_\pi)$ equals 1 if y_π begins with a valid `<think>...</think>` block and 0 otherwise.

For image-to-CAD, the input images do not contain absolute scale information. Therefore, we normalize the solid geometry before computing the reward:

$$n(\Omega) = \left\{ \frac{\mathbf{x} - \bar{\mathbf{x}}}{\sqrt{\frac{\text{tr}(\mathbf{I})}{2 \times \text{Vol}(\Omega_2)}}} \mid \mathbf{x} \in \Omega \right\}, \quad (14)$$

where \mathbf{I} is the inertia matrix and $\bar{\mathbf{x}}$ is the centroid of Ω .

Experiment

We present the experimental details and evaluate our method on text-to-CAD and image-to-CAD generation tasks, com-

Methods	DeepCAD				Fusion 360 (Out-of-Distribution)			
	P -F1 \uparrow	Median CD \downarrow	Mean CD \downarrow	IR \downarrow	P -F1 \uparrow	Median CD \downarrow	Mean CD \downarrow	IR \downarrow
GPT-4o	50.55	107.55	165.67	15.14	55.10	125.68	187.90	19.81
DeepSeek-V3	50.41	110.67	167.90	15.43	<u>54.34</u>	135.10	190.27	18.31
Qwen3-235B-A22B	52.45	109.51	157.49	24.84	<u>53.07</u>	121.83	182.74	26.55
CAD-Translator (Li et al. 2024)	51.28	92.70	162.47	3.87	45.55	192.51	243.59	3.86
Text2CAD (Khan et al. 2024b)	52.35	91.26	160.49	2.75	45.27	254.95	285.95	3.93
CAD-LLaMA (Li et al. 2025a)	<u>60.02</u>	<u>41.77</u>	<u>98.12</u>	0.39	50.47	<u>60.36</u>	<u>142.48</u>	<u>1.29</u>
ReCAD-VL (Ours)	61.48	34.31	72.47	<u>0.81</u>	55.25	34.67	84.89	0.93

Table 1: Text-to-CAD generation results on DeepCAD (left) and Fusion 360 (right). We report average F1 score for primitives (P -F1), mean and median CD and IR. CD is multiplied by 10^3 . **Bold** and underline indicate the best and the second best result.

paring it with SOTA methods to demonstrate its effectiveness.

Experimental Setups

Datasets. Our model, trained on the DeepCAD training set, is evaluated on the DeepCAD and Fusion360 Gallery datasets. After filtering out trivial models, such as cylinders and cubes, we have approximately 90,000 models from the DeepCAD dataset. We extract hierarchical primitives $P \in \mathcal{P}$ from the training set, render images, and calculate pairwise similarities using DINOv2 embeddings. Primitives with similarity scores above 0.95 are deemed duplicates and removed. For fair evaluation, we use Text2CAD (Khan et al. 2024b) descriptions for DeepCAD’s test set and generate descriptions for Fusion360 as per CAD-LLaMA (Li et al. 2025a).

Implementation Details. We select Qwen2.5-VL-7B-Instruct (Bai et al. 2025) as the base model for SFT. The learning rate is set to 1.0×10^{-5} for supervised fine-tuning (SFT) with a context length of 3072, and 2.0×10^{-6} for reinforcement learning (RL) with a context length of 8192, with cosine learning rate scheduling. We employ GPT-4o (Hurst et al. 2024) for both code rewriting and text description generation. Filtering thresholds are set as $\tau_s = 0.95$ and $\tau_h = 0.8$. The weighting coefficients λ_1 and λ_2 are empirically set to 0.1 and 0.9, respectively. We adopt DeepSpeed (Rasley et al. 2020) and Flash-Attention (Dao et al. 2022) to accelerate training. Our framework is built upon veRL (Sheng et al. 2024), where we sample 8 candidate solutions per question q during rollout, with all experiments run on 8 NVIDIA A800 80GB GPUs.

Metrics. We use the Chamfer Distance (CD) to evaluate the geometric alignment between the predicted and ground-truth CAD models, and the Invalidity Ratio (IR) to measure the proportion of invalid CAD models across all tasks. In addition, for text-to-CAD generation, we adopt the average F1 score of primitives (P -F1) (Khan et al. 2024b), while for image-to-CAD generation, we use IOU_{best} from CAD-Coder (Doris et al. 2025), which measures the intersection-over-union under the best alignment between the generated and ground-truth CAD models.




Figure 4: Qualitative comparison of different methods on text-to-CAD and image-to-CAD tasks.

Results on Text-to-CAD Generation

In the text-to-CAD task, we provide the functional interfaces for strong PLMs such as DeepSeek-V3 and GPT-4o, and employ a two-shot approach, with detailed prompts provided in the Appendix. As shown in Table 1, ReCAD-VL achieves more accurate geometry generation compared to previous SOTA methods. On the in-distribution setting, our method outperforms CAD-LLaMA by ~ 7 and 26 in terms of median and mean Chamfer Distance (CD), respectively. This indicates that our approach produces more precise parameters, thereby generating higher-fidelity models and significantly surpassing other baseline methods, as illustrated in Figure 4. These results suggest that both previous methods and PLMs, despite being able to infer correct command types (as reflected by P -F1), struggle to predict accurate parameters.

In the Out-of-distribution setting, our method significantly outperforms all baselines, surpassing CAD-LLaMA by ~ 26 and 58 in median and mean CD, respectively. The comparable performance of our model on both in-distribution and out-of-distribution data highlights its strong generalization ability. Due to the lack of relevant prior knowledge and the absence of training, PLMs, while maintaining consistent performance across distributions, still

Methods	DeepCAD				Fusion 360 (Out-of-Distribution)			
	IOU _{best} ↑	Median CD↓	Mean CD↓	IR↓	IOU _{best} ↑	Median CD↓	Mean CD↓	IR↓
GPT-4o	37.88	198.34	274.23	11.51	41.75	184.45	308.14	8.72
Gemini-2.5-Pro	45.19	160.26	219.98	37.38	44.95	158.46	326.31	39.36
Qwen-2.5-72B	34.51	180.12	258.02	5.01	31.19	175.27	313.75	9.41
CAD-Coder (Doris et al. 2025)	<u>61.23</u>	<u>8.09</u>	<u>73.47</u>	1.05	<u>45.32</u>	<u>84.02</u>	<u>272.06</u>	<u>2.23</u>
ReCAD-VL (Ours)	63.14	7.45	29.61	<u>1.12</u>	54.93	17.01	80.23	0.91

Table 2: Image-to-CAD generation results on DeepCAD (left) and Fusion 360 (right). We report IOU_{best}, mean and median CD and IR. CD is multiplied by 10³. **Bold** and underline indicate the best and the second best result.

struggle to generate precise models. In contrast, previous methods show substantial performance drops on out-of-distribution data, indicating potential overfitting to in-distribution data. This may elucidate their higher invalidity ratio (IR) observed within the Fusion 360 dataset, as exemplified by CAD-LLaMA, which conversely achieves marginally superior results compared to our method in the in-distribution context. Notably, without task-specific training, our model shows impressive performance on CAD-related tasks such as editing and debugging (Figure 3), demonstrating the effectiveness of leveraging PLMs’ intrinsic code generation and reasoning abilities.

Results on Image-to-CAD Generation

In the image-to-CAD task, we use a single image as input and prompt the baseline PLMs using a method similar to that used in the text-to-CAD task. As shown in Table 2, our model outperforms previous methods on nearly all metrics. In particular, we achieve a significant lead over CAD-Coder (Doris et al. 2025) 44 in terms of the mean Chamfer Distance (CD), indicating that our approach is more robust and consistently produces high-fidelity results across different input images. In the Out-of-distribution setting, our model also exhibits strong robustness and generalization ability, outperforming CAD-Coder (Doris et al. 2025) and PLMs by a large margin across all metrics. CAD-Coder (Doris et al. 2025), which is similar to (Guan et al. 2025), converts the code into the CADQuery format. However, these methods neither leverage the pre-existing knowledge and generative capabilities of PLMs nor exploit the inherent strengths of CAD-Query itself. As a result, similar to the text-to-CAD task, their performance drops significantly on out-of-distribution data.

Methods	P-F1	Median CD↓	Mean CD↓	IR↓
SFT <i>only</i>	53.53	84.78	155.67	3.21
RL <i>only</i>	55.61	107.32	179.50	4.77
<i>w/o HPL</i>	59.63	44.64	90.83	2.42
<i>w/o Guidance</i>	60.03	42.85	87.34	0.93
Ours	61.48	34.31	72.47	0.81

Table 3: Ablation experiment on the DeepCAD dataset for the text-to-CAD task.

Ablation Study

We conducted ablation studies on the text-to-CAD task to evaluate the effectiveness of different training strategies in CAD generation. As shown in Table 3, using SFT alone or RL alone results in suboptimal performance. SFT primarily relies on memory-based imitation, which lacks sufficient generalization and often requires repeated sampling to achieve desirable outputs. RL, on the other hand, suffers from low learning efficiency and is limited by the model’s intrinsic capabilities. This suggests that neither method alone is adequate for the CAD domain, particularly in the absence of prior task-specific knowledge. Combining SFT and RL helps mitigate these limitations: SFT introduces external knowledge that the model cannot acquire autonomously, while RL further enhances the model’s ability to generalize by reinforcing the knowledge learned during SFT. Moreover, providing guidance as complementary knowledge enables compensation for the limitations of reinforcement learning (RL) while leveraging the model’s in-context learning capability, thereby further enhancing the generation quality. Without HPL, reconstruction errors and failure rates increase, indicating that it enables gradual compositional learning and yields more reliable CAD structures.

Conclusion

We present ReCAD, a reinforcement learning framework leveraging PLMs to generate precise parametric CAD models. Our method consists of two stages: initially, a supervised fine-tuning phase is employed to transform hardcoded Computer-Aided Design (CAD) scripts into parameterized code, thereby facilitating the generation of precise descriptions. Subsequently, a reinforcement learning phase is implemented, utilizing the parameterized code as a guiding framework. This phase incorporates a hierarchical primitive learning paradigm aimed at developing compositional modeling competencies under a unified reward structure, thus ensuring both geometric and semantic fidelity. Experimental results show that ReCAD performs strongly on in-distribution data and generalizes well to out-of-distribution cases. This demonstrates that by exploiting pre-existing knowledge within PLMs, ReCAD enables precise parameter control while preserving design intent in the generation of CAD models, making it well suited for real-world CAD applications.

Acknowledgments

We thank the anonymous reviewers for their valuable comments. The computations in this research were performed using the CFFF platform of Fudan University.

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, 41–48.
- CadQuery. 2025. CADQuery. <https://github.com/CadQuery/cadquery>. Accessed: 2025-06-01.
- Casey, E.; Zhang, T.; Ishida, S.; Thompson, J. R.; Khasahmadi, A.; Lambourne, J. G.; Jayaraman, P. K.; and Willis, K. D. 2025. Aligning Constraint Generation with Design Intent in Parametric CAD. *arXiv preprint arXiv:2504.13178*.
- Chen, C.; Wei, J.; Chen, T.; Zhang, C.; Yang, X.; Zhang, S.; Yang, B.; Foo, C.-S.; Lin, G.; Huang, Q.; et al. 2025. Cad-crafter: Generating computer-aided design models from unconstrained images. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 11073–11082.
- Cherreddy, S.; and Femiani, J. 2025. SketchDNN: Joint Continuous-Discrete Diffusion for CAD Sketch Generation. *arXiv preprint arXiv:2507.11579*.
- Cherng, J. G.; Shao, X.-Y.; Chen, Y.; and Sferro, P. R. 1998. Feature-based part modeling and process planning for rapid response manufacturing. *Computers & industrial engineering*, 34(2): 515–530.
- Daareyni, A.; Martikkala, A.; Mokhtarian, H.; and Ituarte, I. F. 2025. Generative AI meets CAD: enhancing engineering design to manufacturing processes with large language models. *The International Journal of Advanced Manufacturing Technology*, 1–10.
- Dao, T.; Fu, D.; Ermon, S.; Rudra, A.; and Ré, C. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in neural information processing systems*, 35: 16344–16359.
- Deng, H.; Zou, D.; Ma, R.; Luo, H.; Cao, Y.; and Kang, Y. 2025. Boosting the generalization and reasoning of vision language models with curriculum reinforcement learning. *arXiv preprint arXiv:2503.07065*.
- Dong, Q.; Li, L.; Dai, D.; Zheng, C.; Ma, J.; Li, R.; Xia, H.; Xu, J.; Wu, Z.; Liu, T.; et al. 2022. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*.
- Doris, A. C.; Alam, M. F.; Nobari, A. H.; and Ahmed, F. 2025. CAD-Coder: An Open-Source Vision-Language Model for Computer-Aided Design Code Generation. *arXiv preprint arXiv:2505.14646*.
- Elman, J. L. 1993. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1): 71–99.
- Gao, S.; Wen, X.-C.; Gao, C.; Wang, W.; Zhang, H.; and Lyu, M. R. 2023. What makes good in-context demonstrations for code intelligence tasks with llms? In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 761–773. IEEE.
- Guan, Y.; Wang, X.; Xing, X.; Zhang, J.; Xu, D.; and Yu, Q. 2025. CAD-Coder: Text-to-CAD Generation with Chain-of-Thought and Geometric Reward. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Huang, K.-H.; Qin, C.; Qiu, H.; Laban, P.; Joty, S.; Xiong, C.; and Wu, C.-S. 2025. Why vision language models struggle with visual arithmetic? towards enhanced chart and geometry understanding. *arXiv preprint arXiv:2502.11492*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Khan, M. S.; Dupont, E.; Ali, S. A.; Cherenkova, K.; Kacem, A.; and Aouada, D. 2024a. Cad-signet: Cad language inference from point clouds using layer-wise sketch instance guided attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4713–4722.
- Khan, M. S.; Sinha, S.; Uddin, T.; Stricker, D.; Ali, S. A.; and Afzal, M. Z. 2024b. Text2cad: Generating sequential cad designs from beginner-to-expert level text prompts. *Advances in Neural Information Processing Systems*, 37: 7552–7579.
- Lambert, N.; Morrison, J.; Pyatkin, V.; Huang, S.; Ivison, H.; Brahman, F.; Miranda, L. J. V.; Liu, A.; Dziri, N.; Lyu, S.; et al. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Li, J.; Ma, W.; Li, X.; Lou, Y.; Zhou, G.; and Zhou, X. 2025a. CAD-Llama: leveraging large language models for computer-aided design parametric 3D model generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18563–18573.
- Li, P.; Zhang, W.; Guo, J.; Chen, J.; and Yan, D.-M. 2025b. Revisiting cad model generation by learning raster sketch. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 4869–4877.
- Li, X.; Song, Y.; Lou, Y.; and Zhou, X. 2024. Cad translator: An effective drive for text to 3d parametric computer-aided design generative modeling. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 8461–8470.
- Li, Y.; Lin, C.; Liu, Y.; Long, X.; Zhang, C.; Wang, N.; Li, X.; Wang, W.; and Guo, X. 2025c. CADDreamer: CAD Object Generation from Single-view Images. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 21448–21457.
- Luo, M.; Tan, S.; Huang, R.; Patel, A.; Ariyak, A.; Wu, Q.; Shi, X.; Xin, R.; Cai, C.; Weber, M.; Zhang, C.; Li, L. E.; Popa, R. A.; and Stoica, I. 2025. DeepCoder: A Fully Open-Source 14B Coder at O3-mini Level. <https://www.together.ai/blog/deepcoder>. Accessed: 2025-06-11.
- Ma, W.; Chen, S.; Lou, Y.; Li, X.; and Zhou, X. 2024. Draw Step by Step: Reconstructing CAD Construction Sequences

- from Point Clouds via Multimodal Diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 27154–27163.
- Martin, D. 2023. What is design intent?, 2023. Accessed: January, 19: 2.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.
- Rasley, J.; Rajbhandari, S.; Ruwase, O.; and He, Y. 2020. DeepSpeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 3505–3506.
- Robertson, D.; and Allen, T. J. 2002. CAD system use and engineering performance. *IEEE Transactions on Engineering Management*, 40(3): 274–282.
- Rukhovich, D.; Dupont, E.; Mallis, D.; Cherenkova, K.; Kacem, A.; and Aouada, D. 2025. Cad-recode: Reverse engineering cad code from point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9801–9811.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Seff, A.; Zhou, W.; Richardson, N.; and Adams, R. P. 2021. Vitruvion: A generative model of parametric cad sketches. *arXiv preprint arXiv:2109.14124*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, H.; Liu, P.; Li, J.; Fang, C.; Ma, Y.; Liao, J.; Shen, Q.; Zhang, Z.; Zhao, K.; Zhang, Q.; et al. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Sheng, G.; Zhang, C.; Ye, Z.; Wu, X.; Zhang, W.; Zhang, R.; Peng, Y.; Lin, H.; and Wu, C. 2024. HybridFlow: A Flexible and Efficient RLHF Framework. *arXiv preprint arXiv:2409.19256*.
- Wang, R.; Yuan, Y.; Sun, S.; and Bian, J. 2025a. Text-to-cad generation through infusing visual feedback in large language models. *arXiv preprint arXiv:2501.19054*.
- Wang, S.; Chen, C.; Le, X.; Xu, Q.; Xu, L.; Zhang, Y.; and Yang, J. 2025b. Cad-gpt: Synthesising cad construction sequence with spatial reasoning-enhanced multimodal llms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 7880–7888.
- Willis, K. D.; Pu, Y.; Luo, J.; Chu, H.; Du, T.; Lambourne, J. G.; Solar-Lezama, A.; and Matusik, W. 2021. Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. *ACM Transactions on Graphics (TOG)*, 40(4): 1–24.
- Wong, M.-F.; Guo, S.; Hang, C.-N.; Ho, S.-W.; and Tan, C.-W. 2023. Natural language generation and understanding of big code for AI-assisted programming: A review. *Entropy*, 25(6): 888.
- Wu, R.; Xiao, C.; and Zheng, C. 2021. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6772–6782.
- Wu, S.; Khasahmadi, A. H.; Katz, M.; Jayaraman, P. K.; Pu, Y.; Willis, K.; and Liu, B. 2024. Cadvlm: Bridging language and vision in the generation of parametric cad sketches. In *European Conference on Computer Vision*, 368–384. Springer.
- Xu, J.; Zhao, Z.; Wang, C.; Liu, W.; Ma, Y.; and Gao, S. 2024. Cad-mllm: Unifying multimodality-conditioned cad generation with mllm. *arXiv preprint arXiv:2411.04954*.
- Xu, X.; Jayaraman, P. K.; Lambourne, J. G.; Willis, K. D.; and Furukawa, Y. 2023. Hierarchical neural coding for controllable cad model generation. *arXiv preprint arXiv:2307.00149*.
- Xu, X.; Willis, K. D.; Lambourne, J. G.; Cheng, C.-Y.; Jayaraman, P. K.; and Furukawa, Y. 2022. Skexgen: Autoregressive generation of cad construction sequences with disentangled codebooks. *arXiv preprint arXiv:2207.04632*.
- Yan, J.; Li, Y.; Hu, Z.; Wang, Z.; Cui, G.; Qu, X.; Cheng, Y.; and Zhang, Y. 2025. Learning to reason under off-policy guidance. *arXiv preprint arXiv:2504.14945*.
- Yuan, Y.; Sun, S.; Liu, Q.; and Bian, J. 2025. Cad-editor: A locate-then-infill framework with automated training data synthesis for text-based cad editing. *arXiv preprint arXiv:2502.03997*.
- Zhan, Y.; Zhu, Y.; Zheng, S.; Zhao, H.; Yang, F.; Tang, M.; and Wang, J. 2025. Vision-r1: Evolving human-free alignment in large vision-language models via vision-guided reinforcement learning. *arXiv preprint arXiv:2503.18013*.
- Zhang, E.; Yan, X.; Lin, W.; Zhang, T.; and Lu, Q. 2025a. Learning Like Humans: Advancing LLM Reasoning Capabilities via Adaptive Difficulty Curriculum Learning and Expert-Guided Self-Reformulation. *arXiv preprint arXiv:2505.08364*.
- Zhang, Z.; Liu, K.; Liu, J.; Wang, W.; Lin, B.; Xie, L.; Shen, C.; and Cai, D. 2025b. GeoCAD: Local Geometry-Controllable CAD Generation. *arXiv preprint arXiv:2506.10337*.
- Zhang, Z.; Sun, S.; Wang, W.; Cai, D.; and Bian, J. 2024. Flexcad: Unified and versatile controllable cad generation with fine-tuned large language models. *arXiv preprint arXiv:2411.05823*.
- Zhao, Y.; Wang, X.; Liu, J.; King, I.; and Huang, Z. 2025. Towards Geometry Problem Solving in the Large Model Era: A Survey. *arXiv:2506.02690*.