

Seek-CAD: A Self-refined Generative Modeling for 3D Parametric CAD Using Local Inference via DeepSeek

Xueyang Li*, Jiahao Li*, Yu Song, Yunzhong Lou, Xiangdong Zhou✉

School of Computer Science, Fudan University, Shanghai, China

{xueyangli21, lijh23, songy23}@m.fudan.edu.cn {yzlou20, xdzhou}@fudan.edu.cn

Abstract

The advent of Computer-Aided Design (CAD) generative modeling will significantly transform the design of industrial products. The recent research endeavor has extended into the realm of Large Language Models (LLMs). In contrast to fine-tuning methods, training-free approaches typically utilize the advanced closed-source LLMs, thereby offering enhanced flexibility and efficiency in the development of AI agents for generating CAD parametric models. However, the substantial cost and limitations of local deployment of the top-tier closed-source LLMs pose challenges in practical applications. The Seek-CAD is the pioneer exploration of locally deployed open-source inference LLM DeepSeek-R1 for CAD parametric model generation with a training-free methodology. This study is the first investigation to incorporate both visual and Chain-of-Thought (CoT) feedback within the self-refinement mechanism for generating CAD models. Specifically, the initial generated parametric CAD model is rendered into a sequence of step-wise perspective images, which are subsequently processed by a Vision Language Model (VLM) alongside the corresponding CoTs derived from DeepSeek-R1 to assess the CAD model generation. Then, the feedback is utilized by DeepSeek-R1 to refine the initial generated model for the next round of generation. Moreover, we present an innovative 3D CAD model dataset structured around the SSR (Sketch, Sketch-based feature, and Refinements) triple design paradigm. This dataset encompasses a wide range of CAD commands, thereby aligning effectively with industrial application requirements and proving suitable for the generation of LLMs. Extensive experiments validate the effectiveness of Seek-CAD under various metrics.

1 Introduction

The CAD parametric model (also called design history), as a crucial role in Computer-Aided Design (CAD), indicates the design logic of 3D CAD models, where its command and parameter sequences can be quickly edited to create or modify the shape of a 3D object [54, 21, 32]. However, constructing a parametric CAD model from scratch is time-consuming and hinders the development of automation in industrial manufacturing. Hence, it attracts much attention on generative CAD modeling recently, including many interesting applications such as CAD parts assembly [20, 52, 53], shape parsing [24, 40], CAD parametric model generation [54, 57, 29], and cross-modal CAD generation (e.g., Point cloud-to-CAD, Text-to-CAD, and etc.) [22, 26, 8, 28, 27, 33, 30, 63, 23]. Under the current trend of LLMs and VLMs performing outstandingly in computer vision tasks [50, 41, 47, 16, 64, 3], integrating LLMs and VLMs in generative CAD modeling will pave the way for future innovations in smart design systems.

Fine-tuning is a commonly used method for adapting general LLMs to domain-specific applications. In contrast to fine-tuning approaches, training-free approaches typically utilize advanced closed-

*Equal Contributions

source LLMs i.e. GPT-4o [1], thereby offering enhanced flexibility and efficiency in the creation of AI agents for generating CAD parametric models [61, 2, 25]. However, the substantial cost and limitations of the local deployment of the top-tier closed-source LLMs pose challenges in practical applications. Inspired by recent DeepSeek-R1 [12] with advanced capability of reasoning, we deploy a DeepSeek-R1-32B-Q4 locally without training or finetuning as the backbone of our approach to explore its capability for generative CAD modeling. To the best of our knowledge, this is the first attempt to explore DeepSeek-R1 on generative CAD modeling.

Specifically, we present Seek-CAD, a training-free generative framework for CAD modeling based on DeepSeek-R1-32B-Q4. By employing a retrieval-augmented generation (RAG) strategy on a local CAD code corpus, Seek-CAD produces Python-like CAD code for CAD parametric model generation. To ensure that the generated code aligns with prompt descriptions and faithfully encodes geometric features (e.g., *fillet*, *chamfer*) and constraints (e.g., *tangency*, *orthogonality*), we introduce a step-wise visual feedback mechanism that guides and refines the modeling process. In particular, we render step-wise perspective images that visually capture each stage of the CAD modeling process. These images are then evaluated using Gemini-2.0 [45] to assess their alignment with the chain-of-thought (CoT) from DeepSeek-R1-32B-Q4. The resulting feedback is incorporated to iteratively refine the initial code and parameters.

Furthermore, we propose a novel CAD design paradigm called SSR (Sketch, Sketch-based feature, and Refinements), where each model is represented as a sequence of SSR triples, each consisting of a sketch, a sketch-based feature (e.g., *extrude* and *revolve*), and optionally, refinement features (e.g., *chamfer*, *fillet* and *shell*). Complex shapes are constructed through boolean operations across SSR units. To support refinement features, we introduce a simple yet effective reference mechanism, termed *CapType* (Figure 3), which establishes explicit links between topological primitives in the sketch and their resulting primitives generated during modeling. To evaluate our approach, we construct a new CAD dataset of 40k samples following the SSR modeling paradigm. The dataset covers diverse CAD features not included in existing datasets, and each sample is paired with a textual description generated by GPT-4o [1]. For further details, please see Section A.2. We plan to release the dataset publicly in the future.

In summary, our key contributions are as follows: (i) We present Seek-CAD, a training-free framework leveraging the locally deployed DeepSeek-R1. It incorporates a self-refinement capability through a sequential visual and CoT feedback mechanism, which enhances the generative modeling of CAD designs and significantly contributes to the effective generation of diverse CAD parametric models. (ii) We present an innovative SSR design paradigm, which serves as an alternative to the conventional SE paradigm and demonstrates enhanced suitability for the generation of complex CAD models. (iii) Experimental results demonstrate that Seek-CAD can generate diverse and complex parametric CAD models with high geometric fidelity, enabling precise parametric control while supporting the generation of diverse CAD models.

2 Related Work

2.1 Generative CAD Modeling

Generative CAD modeling has made a significant step recently [56, 58, 18, 8, 13], notably with transformer-based models, which treat CAD commands as parametric sequential data for learning and generation with a feed-forward or an auto-regressive strategy [54, 56, 57, 10]. Besides, diffusion-based [15, 33] methods have also been adopted to achieve the controllable generation or reconstruction of parametric CAD sequences [33, 62]. Mamba-CAD [29] makes a step forward to handle longer parametric CAD sequences to generate complex CAD models. As a CAD model contains explicit parametric commands, which implicitly indicate its design logic and shape geometry, it can be seen as a kind of multi-modal data inherently, which also raises many applications of generating or reconstructing CAD models from point clouds, images, and texts [27, 60, 8, 22, 28]. These methods mainly focus on the SE (Sketch-Extrusion) paradigm, which supports only a limited set of simple CAD operations, making it incapable of generating diverse and complex CAD models that meet real-world design requirements. Instead, Seek-CAD adopts the SSR paradigm, which enables the inclusion of diverse CAD commands such as *fillet*, *chamfer*, and *shell*, supporting the creation of more complex CAD models, which is closer to industrial requirements. From another

perspective, Seek-CAD is a self-refined framework for generative CAD modeling without any training or finetuning, which is also different from these efforts under a training strategy.

2.2 LLMs For CAD

With the success of LLMs in Computer Vision (CV) tasks [44, 35, 42, 11], many interesting applications of leveraging LLMs for CAD-related tasks have raised much attention recently. Text2CAD [23] generates parametric CAD models from natural language instructions using a transformer-based network [46] with an annotation pipeline for CAD prompt using Mistral [4] and LLaVA-NeXT [31]. CAD-MLLM [55], a unified system to utilize a LLM to align multimodal inputs with parametric CAD sequences for generating CAD models. CAD-assistant [34] uses a Vision-Large Language Model (VLLM) with tool-augmented planning to iteratively generate and adapt CAD designs via Python API of [9]. FI2CAD [36] introduces a multi-agent system, a LLM-based MAS architecture for CAD development processes that mimics an engineering team to automatically generate and refine parametric CAD models with human feedback. CAD-Llama [25] enables pretrained LLMs to generate parametric 3D CAD models through the adaptive pretraining on Structured Parametric CAD Code (SPCC). Unlike these methods, our approach mainly focuses on a training-free strategy with the self-refined capability for generative CAD modeling.

The most related to our Seek-CAD are other two training-free frameworks, 3D-PreMise [61] and CADCodeVerify [2]. 3D-Premise enhances CAD code refinement by supplying GPT-4 [1] with an image of a whole object and its initial description, allowing it to identify and correct discrepancies between the intended design and the generated CAD code. Unlike 3D-Premise, CADCodeVerify prompts GPT-4 [1] to generate and answer a set of questions based on the initially provided description of a 3D CAD object, to adjust any discrepancies between the generated CAD model and the description. Compared to them, Seek-CAD differs in the following aspects: (i) Seek-CAD uses DeepSeek-R1-32B-Q4, an open-source reasoning model for local deployment, integrated with a CAD code corpus via a retrieval-augmented generation framework, instead of relying on proprietary LLMs like GPT-4o. (ii) Previous research evaluates object renderings in VLMs by focusing only on final CAD forms, overlooking intermediate phases, which limits feedback for complex models. Seek-CAD addresses this with a step-wise feedback mechanism, showing both final and intermediate shapes to enhance VLM feedback. (iii) Previous methods use predefined question templates to evaluate VLM alignment with descriptions and images. Seek-CAD, however, guides VLMs to assess alignment between DeepSeek-R1’s chain-of-thought (CoT) and step-wise images, as CoT effectively illustrates the design logic, enabling clearer VLM understanding (Sec 3.2).

3 Seek-CAD Framework

In this section, we introduce the Seek-CAD pipeline, which integrates a local inference process (Sec 3.1) and a step-wise visual feedback strategy (Sec 3.2) to generate and progressively refine CAD models, based on visual alignment signals derived from step-wise renderings. The framework of Seek-CAD is illustrated in Figure 1.

3.1 Local Inference Pipeline

LLMs are widely applied across various domains through training or finetuning, both of which require substantial computational resources. In contrast, performing local inference with LLMs can significantly reduce the requirement for computational resources. Inspired by the advancements in reasoning capabilities of DeepSeek-R1, we make a step forward to explore its potential for CAD generative modeling without additional training or finetuning.

(1) Pipeline Definition. This pipeline takes the text T as the input. It is a statement of how to design a CAD model step by step, or it can also be a description of the geometric appearance of objects. The goal can be defined as: $I_0 \sim P(I_0 | T; H)$ that maps the input text T to an initial set of CAD code I_0 represented in the SSR paradigm (Sec 4), where H denotes our Seek-CAD.

(2) Knowledge Constraint. Similar to other LLMs, DeepSeek-R1 may exhibit hallucinations [43, 17, 19], occasionally generating CAD parametric models that deviate from the SSR paradigm. To address it, we propose the knowledge constraint $Cons = (\Phi, \mathcal{D}, \mathcal{E})$ as the system prompt to make DeepSeek-R1 generate CAD code following the SSR paradigm, with the constraint $Cons$ consisting of three parts: Φ specifies its functionality, $\mathcal{D}$ documents the SSR schema, and $\mathcal{E}$ provides an example




Figure 1: The overview of our Seek-CAD framework. The whole pipeline can be divided into two parts consisting of "Initial CAD Code Generation" and "CAD Code Refinement", which are both embedded with a knowledge constraint depicted in Sec. 3.1 to guide DeepSeek-R1 to generate CAD code following the SSR paradigm (Sec. 4). For the first part, a given query T is enhanced by conducting RAG on a local CAD corpus that consisting 10,000 CAD models. Next, Top-3 retrieved candidates would be concatenated with T to trigger DeepSeek-R1 to generate an initial CAD code I_0 . For the second part, I_0 would go through the Step-wise Visual Feedback with CoT to have the iteration refinement. To achieve this, we first utilize a rendering script $R(*)$ to obtain step-wise images of I_0 , which can represent the intermediate and ultimate shape of the object (M_I , M_U) simultaneously. $\oplus$ denotes the concatenation of SSR triplets, where each triplet, represented by S_i (Sec. 4), is rendered along with all its preceding triplets to preserve the correlations between object entities. (More details in Sec. 3.2). Next, the step-wise images are fed into Gemini-2.0 to assess their alignment with the CoT from DeepSeek-R1. This feedback determines whether the current code I_k is reasonable. In practice, we set $k = 2$ as the maximize iterations of code refinement.

pairing a textual description with SSR-based CAD code, as shown in Figure 6 of Appendix A.1.

(3) Retrieval Augment Generation (RAG). In order to augment the recognition of Seek-CAD within the SSR paradigm, we construct a SSR-based CAD corpus $\mathcal{C}_{SSR} = \{(d_i, c_i)\}_{i=1}^N$ to equip Seek-CAD with the capability to perform RAG. d_i is the textual description and c_i its corresponding CAD code, with $N = 10,000$ CAD models. As usual settings of conducting RAG, we select a hybrid search combining a vector-based and full-text strategy. Specifically, for a given T , the similarity score is computed by:

$$g_i^{\text{final}} = \lambda \cdot g_i^{\text{vec}} + (1 - \lambda) \cdot g_i^{\text{full}}, \quad \lambda \in [0, 1], \quad (1)$$

where $g_i^{\text{vec}} = \cos(z_T, z_i)$, $z_T = \mathbf{e}_{\text{vec}}(T)$, $z_i = \mathbf{e}_{\text{vec}}(d_i)$. $\mathbf{e}_{\text{vec}}$ denotes an embedding model, bge-m3 [5]. $g_i^{\text{full}} = \mathbf{U}_{\text{index}}(T, d_i)$ is calculated with a traditional inverted index without vectorization. The top- k candidates are selected as $\mathcal{R}_T = \text{TopK}(g_i^{\text{final}})$. Practically, we set $\lambda = 0.3$, top- $k = 3$.

(4) Initial CAD Code Generation. To incorporate the retrieved content into the input context, we make each retrieved pair $(d_j, c_j) \in \mathcal{R}_T$ concatenated with T . By integrating the knowledge constraints Cons , the initial CAD code I_0 can be obtained as follows:

$$I_0 \sim P(I_0 \mid T \oplus (d_j, c_j), \text{Cons}). \quad (2)$$

The initially generated CAD code I_0 is occasionally subject to compilation failures with the geometry kernel (e.g., PythonOCC [39] utilized in this study), which are attributed to syntax errors E . For remediation, a pattern template Q is employed to automatically rectify E in I_0 , addressing issues such as mismatched parentheses and incorrect capitalization of variable names, depicted as: $I_0 \sim P(I_0 \mid Q)$.

(5) CAD Code Refinement. Upon addressing the aforementioned syntax errors, we employ the geometry kernel to directly render I_0 , thereby acquiring its sequential perspective images (Sec 3.2). Subsequently, the sequential perspective images, in conjunction with the CoT obtained from DeepSeek-R1, are input into Gemini-2.0 for the purpose of evaluating their alignment. Finally, the step-wise visual feedback (Sec 3.2), F_{call} , would be delivered back to DeepSeek-R1 again to refine I_0 to get final CAD code I . This process may iterate N times based on whether F_{call} is positive or negative. To make it clear, we denote F as an indicator to represent the statement of F_{call} . For example, if $L = 1$, it means F_{call} is positive, that is, the CoT and step-wise perspective images are well matched and do not require any further modification, which can be defined with:

$$I = \begin{cases} \{I_m\}_{m=0}^k & L = 1 \\ I_k \sim \{P(I_k \mid I_{k-1}, F_{\text{call}}, \text{Cons})\}_{k=1}^N & L = 0 \end{cases}. \quad (3)$$

Analogous to the initial CAD code I_0 , which bears the potential for syntax errors, a syntax check shall also be executed on I_k during each iterative step, as delineated in: $I_k \sim P(I_k \mid Q)$. In practice,

the maximum iteration step $N = 2$ is established in our work to avert superfluous adjustments, which may arise from the hallucinations of Gemini-2.0. Further elaboration on the iteration step within the refinement stage can be found in the "Refinement" section of Sec 5.2.

3.2 Step-wise Visual Feedback with CoT for Refinement

A significant capability of reasoning LLMs lies in their ability to refine and optimize initial outputs in response to feedback. The self-refined strategy has been adopted in CAD parametric model generation [61, 2]. As CAD command sequences can be easily rendered into the perspective image of an object by using geometry kernel tools like PythonOCC, this can shift the aligned judgment from "Command-Description" to "Image-Description", which can be well handled by VLMs. In this study, we propose a novel step-wise visual feedback (SVF) strategy, which leverages not only the imagery of the object's final form but also retains visuals depicting the intermediate configurations throughout the entire construction process. Additionally, it incorporates the CoTs from DeepSeek-R1 to guide the VLM during its evaluation.

(1) Obtain Step-wise Images. Specifically, the initial generated CAD code I_0 is converted into a sequence (S_{seq}) of n SSR triplets, S_i (Recall Sec 4), as defined with $I_0 \rightarrow S_{seq} = [S_1, S_2, \dots, S_n]$. Next, we leverage the rendering script $R(*)$ (based on PythonOCC) to subsequently compile this sequence S to generate their corresponding perspective images. To better illustrate the construction process from a visual perspective, we capture both the intermediate object shapes, M_I , and the ultimate object shape, M_U . To capture M_I , for each step $k \in \{1, \dots, n\}$ in the construction sequence $S_1 \oplus S_2 \oplus \dots \oplus S_k$, the corresponding image in M_I is rendered by $R(*)$. To maintain the correlations between entities of the object and make them visualized clearly, this rendering process would highlight the entity from the current SSR triplet S_k while entities from all prior SSR triplets S_j (where $j < k$) are hidden as $\bar{S}_j$. To achieve this, we adjust the transparency in $R(*)$ to ensure that the current shape of S_k to the S_j intermediate shape is emphasized. To capture the ultimate shape of the object, M_U , the complete sequence S_{seq} would be directly rendered without hiding any entities. The rendering process can be found in the "CAD Code Refinement" of Figure 1. Finally, each I_0 can be defined with a set of perspective images M consisting of M_I and M_U :

$$M_I = [R(S_1), R(\bar{S}_1 \oplus S_2), \dots, R(\bar{S}_1 \oplus \bar{S}_2 \oplus \dots \oplus S_n)], \quad (4)$$

$$M_U = R(S_1 \oplus S_2 \oplus \dots \oplus S_n), \quad (5)$$

$$M = [M_I, M_U], \quad (6)$$

where $\oplus$ represents the concatenation of SSR triplets. In practice, only a single image is rendered for each S_K . Benefiting from the highlight of each entity in the step-wise images, this effectively avoids occlusion issues encountered in the single-view rendering.

(2) Query for Calling Feedback. Instead of using VLMs to judge the alignment between the initial description T and M , Seek-CAD directly make it judge the alignment between the thought from DeepSeek-R1, $CoT = [t_1, t_2, \dots, t_m]$, and M . $t_i \in CoT$ depicts a clear design logic of the current step according to the initial description T , which is highly compatible with current step-wise image S_K . This would help VLMs understand clearly how the object is constructed. Next, we feed M along with its corresponding thought CoT into Gemini (VLM) and prompt it to judge the alignment between M and CoT to generate feedback F_{call} , this process can be defined as:

$$F_{call} \sim P(F_{call} \mid G, M, CoT), \quad (7)$$

where G is our prompt query designed with two rules for generating F_{call} : (i) If M and CoT are well aligned, F_{call} should be a positive feedback to return. (ii) If M and CoT are mismatched, F_{call} would outline a clear statement to point out discrepancies between M and CoT , which would be delivered back to DeepSeek-R1 to again refine the generated CAD code (Recall Equation 3).

4 SSR Triple Design Paradigm

(1) SSR Triple Design definition. The Sketch and Extrude (SE) paradigm persists as the predominant approach in contemporary feature-based CAD modeling [22, 57, 33, 30, 63, 49, 48, 59, 38], attributed to its versatility and facilitation of parametric editing. Current large-scale datasets for parametric sequence modeling [54, 51] also adopt this paradigm. Nevertheless, the command sets provided by preceding datasets are constrained to basic operations such as *sketch* and *extrude*, and the curves




Figure 2: The SSR Design Paradigm. Each CAD model is constructed as a sequence of SSR triplets, where each triplet consists of a *sketch*, a sketch-based feature (e.g., *extrude*, *revolve*), and optional refinement features (e.g., *shell*, *chamfer*, *fillet*). Topological primitives are traced using the *CapType* reference system (*START*, *SWEPT*, *END*) during modeling operations. Final shapes are formed by applying boolean operations (e.g., Union, Cut, Intersect) between the outputs of SSR triplets.

incorporated within sketches are predominantly confined to elementary types. Consequently, contemporary research endeavours leveraging these datasets tend to yield simpler and less diverse geometric forms, failing to adequately address the complexities inherent in real-world design requirements. To address this, we introduce a novel modeling paradigm called SSR (Sketch, Sketch-based feature, and Refinements), in which each modeling step is represented as a SSR triplet S :

$$S = (s, f, \langle r_1, r_2, \dots, r_k \rangle \text{ or } \emptyset), \quad (8)$$

where $k \geq 0$, s denotes a 2D *sketch* feature, $f \in \mathcal{F}$ is a sketch-based feature such as *extrude* or *revolve*, and $\langle r_1, r_2, \dots, r_k \rangle$ is an ordered sequence of zero or more refinement features, where each $r_i \in \mathcal{R}$ and $\mathcal{R}$ includes features such as *fillet*, *chamfer*, and *shell*. Each SSR triplet S_i is compiled into a 3D geometry $\mathcal{S}_i = \text{CAD_Kernel}(S_i)$ via a CAD kernel like PythonOCC, and a complete CAD model $\mathcal{M}$, composed of n SSR triples, is represented as a sequence of geometry units joined by boolean operations:

$$\mathcal{M} = \langle \mathcal{S}_1, \text{op}_1, \mathcal{S}_2, \text{op}_2, \dots, \mathcal{S}_n \rangle, \quad (9)$$

where $\text{op}_i \in \{\text{Union}, \text{Cut}, \text{Intersect}\}$ denotes a boolean operation applied between adjacent geometry units. An example of this process is shown in Figure 2.

(2) CapType Reference Mechanism. In the context of complex geometric modeling, the integration of refinement features requires referencing particular topological primitives generated during intermediate modeling phases, which are not retained within the design history (parametric model). Consequently, we introduce the *CapType* reference mechanism to address this issue.

Given an SSR triplet $S = (s, f, \langle r_1, r_2, \dots, r_k \rangle)$, where s is the 2D *sketch* containing a set of primitives $\mathcal{A} = \{a_1, a_2, \dots, a_m\}$, we define the intermediate geometry $\mathcal{S}' = \text{CAD_Kernel}(S')$, where $S' = (s, f)$ omits the refinement operations. The resulting geometry $\mathcal{S}'$ contains a set of primitives $\mathcal{B} = \{b_1, b_2, \dots, b_n\}$. The *CapType* reference mechanism defines a mapping $\mathcal{A} \rightarrow \mathcal{B}$ as:

$$\phi(a, C) \rightarrow b, \quad a \in \mathcal{A}, \quad b \in \mathcal{B}, \quad C \in \{\text{START}, \text{END}, \text{SWEPT}\}, \quad (10)$$

where C denotes the *CapType* category: *START* and *END* refer to the primitives at the start and end of the 3D operation f , respectively, while *SWEPT* refers to the primitives generated along the trajectory of the operation, as illustrated in Figure 3. This mechanism allows each refinement operation $r_i \in \mathcal{R}$ to reference a specific primitive $b \in \mathcal{B}$ via $\phi(a, C)$, enabling precise and reliable identification. As illustrated in Figure 5 and supplementary Figure 7, the method enables the generation of complex and diverse CAD models that align well with real-world design requirements.




Figure 3: Illustration of the proposed *CapType* reference mechanism.

5 Experiments

In this section, we perform evaluations of the performance of Seek-CAD and showcase its application across a range of downstream tasks.




Figure 4: (a) Visual illustrations of CAD generative comparison. (b) The visualizations of refinement capability through the SVF strategy (Recall Sec 3.2). Please enlarge to 225% to see the text clearly.

5.1 Experimental Settings

(1) Metrics. To evaluate how precisely the generated CAD commands depict the 3D object, we sample 2000 points separately from CAD models in Ground Truth GT and generated CAD models D . We use **Chamfer Distance (CD)**, **Hausdorff Distance (HD)**, and **Intersection over the Ground Truth (IoGT)** [2] to measure the differences between GT and D . Besides, we utilize **G-Score** (score from 1 to 5, allowing decimals by Gemini-2.0) to judge the alignment between the description and the perspective image of its corresponding generated CAD model. To prove that the generative model is not a mere replication of the local CAD corpus, we also compute a **Novel** metric as: $\frac{1}{n} \sum_{i=1}^n \mathbb{I}[s(I_A, I_{B_i}) < \tau] \geq \rho$, where I_A is a rendering image of the generated CAD model by SeekCAD and I_{B_i} is a rendering image of each CAD model in the local CAD corpus, and $s(*)$ is a similarity function ($\tau = 0.8$, $\rho = 0.8$). In practice, we use ResNet-50 [14] to encode I_A and I_{B_i} before calculating their similarity. Note that cases failed to compile would be excluded when calculating these metrics. Given the generated CAD commands could be failed to compile, we further add **Pass@k** to better understand the mechanism of components in Seek-CAD. **Pass@k** denotes for each description, the probability of at least one set of generated CAD commands being compiled successfully in k times generation. We report mean values under all metrics.

(2) Comparison Methods. We have witnessed two training-free efforts of generative CAD modeling with the self-refined strategy: 3D-PreMise [61] and CADCodeVerify [2]. To compare with them fairly, only the refined strategy in Seek-CAD (SVF) is separately replaced with refined strategies in 3D-PreMise and CADCodeVerify, while the rest of the components in Seek-CAD are all consistent.

(3) Implementation Details. We deploy DeepSeek-R1:32B in Q4 quantization version as a backbone of Seek-CAD on one NVIDIA RTX 3090 GPU with Ollama [37]. The context length is set to 15,000 to endow inference speed of 21.78 tokens per second. For the inference stage, we set temperature $T = 0.7$, top-p as 0.8 to make it capable of generating diverse CAD models in different trials. For RAG settings, we employ Dify [6] and bge-m3 [5], details refer to Appendix A.4. For the refinement stage, we adopt the Gemini-2.0 API [45] from Google AI to judge the discrepancies between step-wise perspective images and their corresponding thoughts (DeepSeek-R1:32B-Q4).

Table 1: The quantitative results of generation ability tested on 500 CAD models.

Method	CD↓	HD↓	IoGT↑	G-Score↑	Novel↑
3D-PreMise	0.2203	0.6137	0.6315	3.2022	49.57%
CADCodeVerify	0.2164	0.5917	0.6562	3.3927	55.38%
Seek-CAD	0.1979	0.5566	0.7226	3.5185	64.04%

Table 2: The quantitative results of comparing refinement rounds tested on 500 CAD models.

Refine Round	Pass@2↑	CD↓	HD↓	IoGT↑	G-Score↑
0	0.77	0.2275	0.6194	0.6183	3.1401
1	0.72	0.1979	0.5566	0.7226	3.5185
2	0.55	0.1966	0.5548	0.7347	3.5314

5.2 Experimental Results

(1) Generation. To avoid data overlapping, we construct a test set consisting of an additional 500 CAD models which are completely different from 10,000 CAD models in the local CAD corpus. Based on our framework, we compare our SVF with two other visual feedback strategies in one-step refinement, as shown in Table 1. As shown in Table 1, the SVF strategy employed in Seek-CAD

significantly outperforms the feedback methods used in 3D-Premise and CADCodeVerify across all metrics. This demonstrates that the CoT and step-wise images utilized in the SVF strategy are indeed well-aligned, and the object design logic contained within the CoT empowers the VLM (Gemini-2.0) to provide clearer feedback, enabling the more precise refinement. Besides, all feedback methods based on our framework achieve close to or over 50% under the Novel metric. This indicates that while the local CAD corpus serves as a constraint on CAD command generation, it does not fully limit the creative ability of our framework to generate novel CAD models. Figure 4(a) gives comparable showcases of Ground Truth *GT* and the generated CAD models from Seek-CAD with different refined strategies. Compared to 3D-PreMise and CADCodeVerify, the SVF embedded in Seek-CAD makes the generated CAD model closer to *GT*, which again proves that SVF (Sec 3.2) is reasonable and effective to integrate step-wise images and the thought from DeepSeek-R1.

(2) Refinement Step. To better understand the impact of visual feedback on CAD model generations, we explore the effects of performing 0, 1, and 2 rounds of SVF within Seek-CAD, which can be found in Table 2. Compared to Round 0, the generated CAD models show significant improvements across four metrics (CD, HD, IoGT, G-Score) after undergoing corrections in Round 1 and Round 2, which proves the effectiveness of SVF. Additionally, we found that the gains diminish significantly with an increasing round of SVF iterations. For instance, compared to Round 1, Round 2 shows only marginal improvements on the four metrics (e.g., IoGT from 0.7226 to 0.7347), but it increases the probability of code compilation failures (Pass@2 from 0.72 to 0.55). There are two main reasons for this: (i) Limited by the inherent reasoning ability of the base model (DeepSeek-R1:32B-Q4) in Seek-CAD, it is difficult for the model to make perfectly precise corrections based on SVF feedback. (ii) There is a certain probability that the generated CAD commands will fail to compile during each generation process. Each additional round of refinement would increase its probability. Hence, we set the maximum iteration step of 2 for the refinement stage to avoid excessive and potentially unnecessary modifications. Two cases are depicted in Figure 4(b). In the right-hand instance, the initially generated CAD model exhibits a reversed extrusion direction, resulting in the failure to form holes in the corners and the middle of the object, which is resolved after one round of refinement. It again demonstrates the effectiveness of the SVF strategy.

(3) Ablation Study. To figure out the mechanism of Seek-CAD, we conduct ablation studies on 200 CAD models by removing some components in Seek-CAD to achieve an additional 7 frameworks tagged from A to G, which separately denote: A (*w/o local CAD corpus*), B (*w/o knowledge constraint*), C (*w/o inter images in SVF*), D (*w/o ultimate image in SVF*), E (*w/o vector search in local CAD corpus*), F (*w/o text search in local CAD corpus*), and G (*w/o CoT in SVF*). For the complete version of Seek-CAD, we denote it as H. Note that all comparisons in the ablation study are based on the one-time refinement. As shown in Table 3, model A completely fails to generate compilable CAD commands, indicating that within a training-free framework, the local CAD corpus indeed imposes constraints to guide model to generate CAD commands based on the SSR format. This is highly effective for the Seek-CAD model, which lacks prior knowledge of SSR, providing a technical approach for quickly achieving CAD model generations without any training. Compared to model H, the impact of models B, E, and F is primarily focused on the compilability of generated CAD commands. In particular, the substantial decline in Pass@1 for model B (from 0.68 to 0.44) highlights the critical role of the system prompt as an effective constraint during CAD command generation. Besides, models E and F also exhibit a significant decrease in four other metrics (e.g., from 0.7451 to 0.6464 in IoGT, 3.8409 to 3.6315 in G-Score). This proves conducting hybrid research in the local CAD corpus is more effective. Furthermore, we found that the image utilization strategy within SVF mainly impacts the precision of the refinement, as evidenced by model C achieving worse scores across all metrics (e.g., 0.2295 vs 0.2114 in CD) compared to model D. This validates that inter-images can provide more visual information than the ultimate-image, thereby enabling the VLM to better understand the construction process of objects and generate clear feedback to enhance further modifications. Model G shows a decline in generation quality (e.g., G-Score from 3.8409 to 3.6120) compared to model H. It indicates CoT helps Gemini-2.0 better understand the step-wise images and improve the quality of feedback.

Table 3: Ablation Studies on 200 CAD models.

	Pass@1↑	Pass@2↑	CD↓	HD↓	IoGT↑	G-Score↑
A	-	-	-	-	-	-
B	0.44	0.64	0.2295	0.6307	0.6287	3.5896
C	0.67	0.79	0.2114	0.5914	0.6713	3.7254
D	0.67	0.80	0.1961	0.5573	0.7036	3.7761
E	0.56	0.77	0.2178	0.5947	0.6563	3.6642
F	0.47	0.68	0.2235	0.6188	0.6464	3.6315
G	0.65	0.76	0.2247	0.6254	0.6373	3.6120
H	0.68	0.81	0.1923	0.5382	0.7451	3.8409




Figure 5: Various Showcases by Seek-CAD. Please enlarge to 180% to see the text clearly.

Finally, the complete model H surpasses other comparable versions across all metrics (e.g., highest score 0.7451 in IoGT), proving the design components in Seek-CAD are all effective and reasonable.

(4) Various Showcases by Seek-CAD. We showcase more capabilities of our Seek-CAD, including **CAD editing**, **Similar generations**, and **Gallery of generated CAD models**. **(i) CAD editing** is an applicative sought in a wide range of industrial sectors. As Seek-CAD is a framework built on DeepSeek-R1, it supports multiple rounds of dialogue to achieve substantial editing that includes but is not limited to "Add", "Remove", and "Scale" modification. Figure 5(a) shows a demo to edit the vanilla generated CAD model based on the user's description iteratively. **(ii) Similar generations.** For the same textual descriptions, Seek-CAD can generate similar CAD models as shown in Figure 5(b), which offers useful suggestions for initial CAD design. Besides, the description in Figure 5(b) contains only functional specifications of the objects, in contrast to that in Figure 5(a) which includes specific dimensional parameters. It indicates Seek-CAD is capable of interpreting functional descriptions to produce valid CAD models without relying on explicit parameter guidance. **(iii) The gallery of generated CAD models.** Figure 5(c) provides some generated CAD models from Seek-CAD with feeding textual descriptions from the test set. This shows that Seek-CAD is capable of generating relatively complex models, which was not demonstrated by previous methods based on SE paradigms or CadQuery.

(5) Limitations and Discussions. Seek-CAD currently has several limitations. First, the step-wise visual feedback contains biases from VLMs due to two main issues: (i) VLMs, confined to their training, struggle to accurately describe object geometries without specific-domain training, and (ii) complex models can't be easily summarized. Second, using DeepSeek-R1 for visual feedback introduces redundant information, as full CoTs are fed to VLMs. Third, DeepSeek-R1-32B-Q4's computational ability for geometric constraints is weaker compared to DeepSeek-R1-671B, needing more refinement for accurate CAD commands. Fourth, the *CapType* referencing system struggles with selecting certain topological primitives, like intersection-generated edges, which need intermediate B-Rep modeling data.

6 Conclusion

We introduce Seek-CAD, a new training-free framework for parametric CAD model generation using the open-source DeepSeek-R1 for cost-efficient inference. It refines CAD commands by integrating visual signals and CoT reasoning with Vision-Language Model evaluation. We also present a corpus of 40,000 annotated CAD models based on the SSR design paradigm, which allows for precise operations like *fillet* and *shell*, which are absent in the traditional Sketch-Extrude paradigm datasets. Seek-CAD surpasses existing training-free methods in generation quality and geometric fidelity, pioneering the use of open-source LLMs in CAD workflows and advancing AI for design automation.

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2, 3
- [2] Kamel Alrashedy, Pradyumna Tambwekar, Zulfiqar Haider Zaidi, Megan Langwasser, Wei Xu, and Matthew Gombolay. Generating cad code with vision-language models for 3d designs. In *The Thirteenth International Conference on Learning Representations*, 2025. 2, 3, 5, 7
- [3] Léore Bensabath, Mathis Petrovich, and Gul Varol. A cross-dataset study for text-based 3d human motion retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1932–1940, 2024. 1
- [4] Devendra Singh Chaplot, Albert q. jiang, alexandre sablayrolles, arthur mensch, chris bamford, devendra singh chaplot, diego de las casas, florian bressand, gianna lengyel, guillaume lample, lucile saulnier, l  lio renard lavaud, marie-anne lachaux, pierre stock, teven le scao, thibaut lavril, thomas wang, timoth  e lacroix, william el sayed. *arXiv preprint arXiv:2310.06825*, 2023. 3
- [5] Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216*, 2024. 4, 7, 17
- [6] Dify. Dify. <https://dify.ai/>. 7, 17
- [7] Docker. Docker. <https://www.docker.com/>. 17
- [8] Elona Dupont, Kseniya Cherenkova, Dimitrios Mallis, Gleb Gusev, Anis Kacem, and Djamila Aouada. Transcad: A hierarchical transformer for cad sequence inference from point clouds. In *European Conference on Computer Vision*, pages 19–36. Springer, 2024. 1, 2
- [9] FreeCAD. Freecad. <https://www.freecad.org/>. 3
- [10] Yaroslav Ganin, Sergey Bartunov, Yujia Li, Ethan Keller, and Stefano Saliceti. Computer-aided design as language. *Advances in Neural Information Processing Systems*, 34:5885–5897, 2021. 2
- [11] Akash Ghosh, Arkadeep Acharya, Raghav Jain, Sriparna Saha, Aman Chadha, and Setu Sinha. Clipsyntel: clip and llm synergy for multimodal question summarization in healthcare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 22031–22039, 2024. 3
- [12] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025. 2
- [13] Haoxiang Guo, Shilin Liu, Hao Pan, Yang Liu, Xin Tong, and Baining Guo. Complexgen: Cad reconstruction by b-rep chain complex generation. *ACM Transactions on Graphics (TOG)*, 41(4):1–18, 2022. 2
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 7
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2
- [16] Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. Vtimellm: Empower llm to grasp video moments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14271–14280, 2024. 1
- [17] Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Transactions on Information Systems*, 43(2): 1–55, 2025. 3
- [18] Pradeep Kumar Jayaraman, Aditya Sanghi, Joseph G Lambourne, Karl DD Willis, Thomas Davies, Hooman Shayani, and Nigel Morris. Uv-net: Learning from boundary representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11703–11712, 2021. 2

- [19] Ling Jiang, Keer Jiang, Xiaoyu Chu, Saaransh Gulati, and Pulkit Garg. Hallucination detection in llm-enriched product listings. In *Proceedings of the Seventh Workshop on e-Commerce and NLP@ LREC-COLING 2024*, pages 29–39, 2024. 3
- [20] Benjamin Jones, Dalton Hildreth, Duowen Chen, Ilya Baran, Vladimir G Kim, and Adriana Schulz. Automate: A dataset and learning approach for automatic mating of cad assemblies. *ACM Transactions on Graphics (TOG)*, 40(6):1–18, 2021. 1
- [21] Benjamin T Jones, Michael Hu, Milin Kodnongbua, Vladimir G Kim, and Adriana Schulz. Self-supervised representation learning for cad. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21327–21336, 2023. 1
- [22] Mohammad Sadil Khan, Elona Dupont, Sk Aziz Ali, Kseniya Cherenkova, Anis Kacem, and Djamila Aouada. 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*, pages 4713–4722, 2024. 1, 2, 5
- [23] Mohammad Sadil Khan, Sankalp Sinha, Talha Uddin, Didier Stricker, Sk Aziz Ali, and Muhammad Zeshan Afzal. Text2cad: Generating sequential cad designs from beginner-to-expert level text prompts. *Advances in Neural Information Processing Systems*, 37:7552–7579, 2024. 1, 3
- [24] Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. Free2cad: Parsing freehand drawings into cad commands. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 1
- [25] Jiahao Li, Weijian Ma, Xueyang Li, Yunzhong Lou, Guichun Zhou, and Xiangdong Zhou. Cad-llama: Leveraging large language models for computer-aided design parametric 3d model generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL <https://arxiv.org/abs/2505.04481>. 2, 3
- [26] Pu Li, Jianwei Guo, Huibin Li, Bedrich Benes, and Dong-Ming Yan. Sfmcad: Unsupervised cad reconstruction by learning sketch-based feature modeling operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4671–4680, June 2024. 1
- [27] Xueyang Li, Haotian Chen, Yunzhong Lou, and Xiangdong Zhou. Cf-cad: A contrastive fusion network for 3d computer-aided design generative modeling. In *International Conference on Database Systems for Advanced Applications*, pages 435–450. Springer, 2024. 1, 2
- [28] Xueyang Li, Yu Song, Yunzhong Lou, and Xiangdong Zhou. 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*, pages 8461–8470, 2024. 1, 2
- [29] Xueyang Li, Yunzhong Lou, Yu Song, and Xiangdong Zhou. Mamba-cad: State space model for 3d computer-aided design generative modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025. 1, 2
- [30] Yuan Li, Cheng Lin, Yuan Liu, Xiaoxiao Long, Chenxu Zhang, Ningna Wang, Xin Li, Wenping Wang, and Xiaohu Guo. Caddreamer: Cad object generation from single-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. 1, 5
- [31] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Lllavanext: Improved reasoning, ocr, and world knowledge, 2024. 3
- [32] Yunzhong Lou, Xueyang Li, Haotian Chen, and Xiangdong Zhou. Brep-bert: Pre-training boundary representation bert with sub-graph node contrastive learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1657–1666, 2023. 1
- [33] Weijian Ma, Shuaiqi Chen, Yunzhong Lou, Xueyang Li, and Xiangdong Zhou. 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*, pages 27154–27163, 2024. 1, 2, 5
- [34] Dimitrios Mallis, Ahmet Serdar Karadeniz, Sebastian Cavada, Danila Rukhovich, Niki Foteinopoulou, Kseniya Cherenkova, Anis Kacem, and Djamila Aouada. Cad-assistant: Tool-augmented vllms as generic cad task solvers? *arXiv preprint arXiv:2412.13810*, 2024. 3
- [35] Muhammad Ferjad Naeem, Muhammad Gul Zain Ali Khan, Yongqin Xian, Muhammad Zeshan Afzal, Didier Stricker, Luc Van Gool, and Federico Tombari. I2mvformer: Large language model generated multi-view document supervision for zero-shot image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15169–15179, 2023. 3

- [36] Felix Ocker, Stefan Menzel, Ahmed Sadik, and Thiago Rios. From idea to cad: A language model-driven multi-agent system for collaborative design. *arXiv preprint arXiv:2503.04417*, 2025. 3
- [37] Ollama. Ollama. <https://ollama.com/>. 7
- [38] Wamiq Reyaz Para, Shariq Farooq Bhat, Paul Guerrero, Tom Kelly, Niloy Mitra, Leonidas Guibas, and Peter Wonka. Sketchgen: Generating constrained cad sketches. 2021. URL <https://arxiv.org/abs/2106.02711>. 5
- [39] Thomas Paviot. "pythonocc". zenodo. <https://doi.org/10.5281/zenodo.3605364>, 2022. 4
- [40] Daxuan Ren, Jianmin Zheng, Jianfei Cai, Jiatong Li, and Junzhe Zhang. Extrudenet: Unsupervised inverse sketch-and-extrude for shape parsing. In *European Conference on Computer Vision*, pages 482–498. Springer, 2022. 1
- [41] Raphael Schumann, Wanrong Zhu, Weixi Feng, Tsu-Jui Fu, Stefan Riezler, and William Yang Wang. Velma: Verbalization embodiment of llm agents for vision and language navigation in street view. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18924–18933, 2024. 1
- [42] Hao Shao, Yuxuan Hu, Letian Wang, Guanglu Song, Steven L. Waslander, Yu Liu, and Hongsheng Li. Lmdrive: Closed-loop end-to-end driving with large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15120–15130, June 2024. 3
- [43] Gaurang Sriramanan, Siddhant Bharti, Vinu Sankar Sadasivan, Shoumik Saha, Priyatham Kattakinda, and Soheil Feizi. Llm-check: Investigating detection of hallucinations in large language models. *Advances in Neural Information Processing Systems*, 37:34188–34216, 2024. 3
- [44] Reuben Tan, Ximeng Sun, Ping Hu, Jui-hsien Wang, Hanieh Deilamsalehy, Bryan A Plummer, Bryan Russell, and Kate Saenko. Koala: Key frame-conditioned long video-llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13581–13591, 2024. 3
- [45] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024. 2, 7
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [47] Junchi Wang and Lei Ke. Llm-seg: Bridging image segmentation and large language model reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1765–1774, 2024. 1
- [48] Ruiyu Wang, Yu Yuan, Shizhao Sun, and Jiang Bian. Text-to-cad generation through infusing visual feedback in large language models. 2025. URL <https://arxiv.org/abs/2501.19054>. 5
- [49] Siyu Wang, Cailian Chen, Xinyi Le, Qimin Xu, Lei Xu, Yanzhou Zhang, and Jie Yang. Cad-gpt: Synthesising cad construction sequence with spatial reasoning-enhanced multimodal llms. 2024. URL <https://arxiv.org/abs/2412.19663>. 5
- [50] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *Advances in Neural Information Processing Systems*, 36:61501–61513, 2023. 1
- [51] Karl DD Willis, Yewen Pu, Jieliang Luo, Hang Chu, Tao Du, Joseph G Lambourne, Armando Solar-Lezama, and Wojciech Matusik. Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. *ACM Transactions on Graphics (TOG)*, 40(4):1–24, 2021. 5
- [52] Karl DD Willis, Pradeep Kumar Jayaraman, Hang Chu, Yunsheng Tian, Yifei Li, Daniele Grandi, Aditya Sanghi, Linh Tran, Joseph G Lambourne, Armando Solar-Lezama, et al. Joinable: Learning bottom-up assembly of parametric cad joints. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15849–15860, 2022. 1
- [53] Ruihai Wu, Chenrui Tie, Yushi Du, Yan Zhao, and Hao Dong. Leveraging se (3) equivariance for learning 3d geometric shape assembly. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14311–14320, 2023. 1
- [54] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6772–6782, 2021. 1, 2, 5, 15

- [55] Jingwei Xu, Chenyu Wang, Zibo Zhao, Wen Liu, Yi Ma, and Shenghua Gao. Cad-mllm: Unifying multimodality-conditioned cad generation with mllm. *arXiv preprint arXiv:2411.04954*, 2024. 3
- [56] Xiang Xu, Karl DD Willis, Joseph G Lambourne, Chin-Yi Cheng, Pradeep Kumar Jayaraman, and Yasutaka Furukawa. Skexgen: Autoregressive generation of cad construction sequences with disentangled codebooks. *arXiv preprint arXiv:2207.04632*, 2022. 2
- [57] Xiang Xu, Pradeep Kumar Jayaraman, Joseph George Lambourne, Karl DD Willis, and Yasutaka Furukawa. Hierarchical neural coding for controllable cad model generation. In *International Conference on Machine Learning*, pages 38443–38461. PMLR, 2023. 1, 2, 5
- [58] Xiang Xu, Joseph Lambourne, Pradeep Jayaraman, Zhengqing Wang, Karl Willis, and Yasutaka Furukawa. Brep-gen: A b-rep generative diffusion model with structured latent geometry. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024. 2
- [59] Mohsen Yavartanoo, Sangmin Hong, Reyhaneh Neshatavar, and Kyoung Mu Lee. Text2cad: Text to 3d cad generation via technical drawings. 2024. URL <https://arxiv.org/abs/2411.06206>. 5
- [60] Yang You, Mikaela Angelina Uy, Jiaqi Han, Rahul Thomas, Haotong Zhang, Suyu You, and Leonidas Guibas. Img2cad: Reverse engineering 3d cad models from images through vlm-assisted conditional factorization. *arXiv preprint arXiv:2408.01437*, 2024. 2
- [61] Zeqing Yuan, Haoxuan Lan, Qiang Zou, and Junbo Zhao. 3d-premise: Can large language models generate 3d shapes with sharp features and parametric control? *arXiv preprint arXiv:2401.06437*, 2024. 2, 3, 5, 7
- [62] Aijia Zhang, Weiqiang Jia, Qiang Zou, Yixiong Feng, Xiaoxiang Wei, and Ye Zhang. Diffusion-cad: Controllable diffusion model for generating computer-aided design models. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 2
- [63] Zhanwei Zhang, Shizhao Sun, Wenxiao Wang, Deng Cai, and Jiang Bian. Flexcad: Unified and versatile controllable cad generation with fine-tuned large language models. In *International Conference on Learning Representations*, 2025. 1, 5
- [64] Yong Zhu, Zhenyu Wen, Xiong Li, Xiufang Shi, Xiang Wu, Hui Dong, and Jiming Chen. Chatnav: Leveraging llm to zero-shot semantic reasoning in object navigation. *IEEE Transactions on Circuits and Systems for Video Technology*, 2024. 1

A Appendix

A.1 Knowledge Constraint

We design a knowledge constraint as the system prompt to guide Seek-CAD in generating CAD code that conforms to the SSR paradigm, which helps reduce hallucination to some extent. It consists of three parts including *State of functionality*, *CAD Code Documentation*, and an *Example* as following:

#Statement of functionality#

You are a role of CAD-assitant to help user create CAD models. You will be given a description of the shape geometry of 3D objects (e.g., triangle, square, rectangle, and prism), and may contain key parameters (e.g., length, height, width, and radius) and some constraints on the entities (e.g., tangent, parallel, and orthogonal). Your job is to help users complete a whole CAD design logic and show the summary of total CAD commands based on their descriptions. You must follow the rule and definition of each CAD command and its parameter from the #CAD Code Documentation# when you give the total CAD commands code. I will give you a #CAD Code d# to show how it works.

#Tips_1# Please give specific and reasonable parameter values to ensure the geometric constraints (e.g., Tangent, Intersect, Orthogonal) between entities when there are multiple sketches.

#Tips_2# You cannot ignore any parameter filled with corresponding CAD commands

#Tips_3# Finally, please use the most effective and simple CAD command to answer. For example, when the user make a query to construct a chamfer edge, please use "Chamfer" command directly instead of using other CAD commands to replace it.

#CAD Code Documentation#

Methods: - 'moveTo(x, y)': Set the starting point of the loop. - 'lineTo(x, y)': Draw a straight line from the current point to '(x, y)'.

- 'threePointArc(p1, p2)': Draw a circular arc from the current point through 'p1' to 'p2'.

- 'circle(radius)': Draw a full circle centered at the current point with the given 'radius'.

- 'splineTo(*points)': Create a smooth spline curve through the given sequence of points from the current point.

- 'pointTag(tag)': Assign a tag to the current point (used after curve method).

- 'curveTag(tag)': Assign a tag to the most recent curve (used after curve method).

Profile Class: Represents a closed shape made up of one or more loops.

Constructor: - 'Profile(tag=None)': Optionally provide a 'tag' to identify the profile.

Method: - 'addLoop(*loops)': Add one or more 'Loop' objects to the profile.

Sketch Class: Defines a 2D sketch on a given plane.

Constructor: - 'Sketch(plane)': Must provide a 'plane' dictionary with:

- 'normal': A 3D vector '[x, y, z]' indicating the plane's normal direction.

- 'origin': A 3D point '[x, y, z]' specifying the sketch's origin.

- 'x': A 3D vector '[x, y, z]' defining the X-axis direction in the plane.

Method: - 'addProfile(*profiles)': Add one or more 'Profile' objects to the sketch.

Shape Creation

'Extrude(sketch, distance=(a, b)) → Shape': Create a 3D solid by extruding the profiles in a sketch.

- 'a': Extrusion distance in the sketch's normal direction. - 'b': Extrusion distance in the opposite direction.

'Revolve(sketch, axis=[point, direction], angle=(a, b)) → Shape': Create a 3D solid by revolving a profile around a given axis.

- 'point': A 3D point '[x, y, z]' on the axis. - 'direction': A 3D vector '[x, y, z]' indicating the axis direction.

- 'a': Angle in the normal direction. - 'b': Angle in the opposite direction.

Shape Class: Represents a solid body. Supports Boolean and modification operations.

Boolean Operations: - 'shape.union(s: Shape) → Shape': Merge the current shape with shape 's'.

- 'shape.cut(s: Shape) → Shape': Subtract shape 's' from the current shape.

- 'shape.intersect(s: Shape) → Shape': Keep only the intersecting volume between the current shape and shape 's'.

Modifiers: - 'shape.Chamfer(width, entities)': Apply a chamfer to tagged edges or curves.

- 'shape.Fillet(radius, entities)': Round corners with a fillet radius. - 'shape.Shell(thickness, entities)': Hollow out the shape with specified wall thickness.

#Example#

Description # The resulting model is an L-shaped prismatic solid with two perpendicular rectangular arms, one extending vertically and the other horizontally, forming the characteristic L shape in 3D space. Additionally, there are four vertical rectangular pillar-like prisms extending downward from each corner of the virtual bounding rectangle.

Generated Answers # Based on your description, I can help you create the CAD command summary for generating the described shape. Here is the complete CAD code command design with valid parameters:

```
sk0 = Sketch(plane={"normal": [0.0, 0.0, 1.0], "origin": [0.0, 0.0, 0.0], "x": [1.0, 0.0, 0.0]})
```

```
p0_0 = Profile()
```

```
p0_0.addLoop(Loop().moveTo(0.0,-2.667).lineTo(0.254,-2.667).lineTo(0.254,-0.254)
```

```
.lineTo(1.8288,-0.254).lineTo(1.8288,0.0)
```

```
.lineTo(0.0,0.0)
```

```
.lineTo(0.0,-2.667))
```

```
sk0.addProfile(p0_0)
```

```
shape0 = Extrude(sk0, distance=0.9144)
```

```
sk1 = Sketch(plane={"normal": [0.0, 0.0, 1.0], "origin": [0.0, 0.0, 0.0], "x": [1.0, 0.0, 0.0]})
```

```
p1_0 = Profile()
```

```
p1_0.addLoop(Loop()
```

```
....
```

```
sk1.addProfile(p1_0,p1_1,p1_2,p1_3)
```

```
shape1 = Extrude(sk1, distance=0.8636)
```

```
shape0 = shape0.union(shape1)
```

Figure 6: The knowledge constraint adopted in Seek-CAD.

A.2 CAD dataset Based on SSR Paradigm

We extend the data parsing methodology of DeepCAD by utilizing Onshape’s Developer API and FeatureScript to parse CAD models based on the ABC dataset. Our method supports core modeling commands including *sketch*, *extrude*, *revolve*, *fillet*, *chamfer*, and *shell*. During the parsing process, models containing unsupported operations (such as mirror) are not immediately discarded. Instead, we filter them only if the proportion of unsupported commands exceeds a threshold of 0.2. Unsupported commands are ignored, and only supported operations are retained for further processing.

For commands like fillet, chamfer, and shell, we first extract the primitives they operate on (e.g., edges for *chamfer* or faces for *shell*), and then use PythonOCC APIs to match these primitives to specific sketch elements through their *CapType* designation. We note that the proposed *CapType* reference system cannot identify certain edges or faces that are not directly associated with sketch-defined primitives such as those generated by the intersection of solid bodies. Therefore, when refinement commands (*chamfer*, *fillet* or *shell*) involve such primitives that cannot be identified via *CapType* schema, we exclude those primitives from the commands. Finally, all parsed data is converted into a unified JSON format compatible with DeepCAD, which can be directly rendered using PythonOCC for visualization.

To further illustrate the characteristics of our dataset, Figure 7 presents a qualitative comparison between our dataset and DeepCAD. The statistical distributions of command sequence lengths and the number of curves per CAD model are shown in Figure 8.




Figure 7: Qualitative comparison between our dataset and the DeepCAD [54] dataset. Compared to DeepCAD, our dataset captures more realistic and structurally complex industrial designs, supporting a broader range of modeling operations such as spline, revolve, chamfer, fillet, and shell. Our extended dataset also features richer geometric details and greater diversity in modeling strategies.




Figure 8: Statistical comparison between our complete dataset and the DeepCAD [54] dataset, which also is a subset of our dataset. The left plot illustrates the distribution of feature lengths, while the right plot shows the distribution of average curve lengths.

A.3 CAD Code Representation

Since existing CAD scripting tools such as CADQuery do not directly support *CapType* reference mechanism and SSR design paradigm, we define a concise and effective Python-based representation tailored for SSR-driven CAD modeling. As shown in Listing 1, the CADShape class encapsulates an SSR triplet and supports boolean operations—such as union, cut, and intersect—with other SSR triplets. The final CAD model is represented by a CADShape instance, which can either be a single SSR triplet or the result of boolean operations among multiple SSR triplets.

For refinement features (e.g., chamfer, fillet, and shell), the target edge or face primitives are determined using the *CapType* introduced in Section 4. In particular, applying a chamfer or fillet to a face affects all edges on that face. Finally, we convert the code into the JSON format similar to DeepCAD, which is then rendered using PythonOCC.

Listing 1: Python interface for SSR-based CAD modeling. The CADShape class provides a concise and extensible representation that supports structured modeling under the SSR design paradigm.

```
class Loop:
    def __init__(self) -> None: ...
    # Define start point of this loop
    def moveTo(self, x: float, y: float) -> Self: ...
    # Draw straight line to point
    def lineTo(self, x: float, y: float) -> Self: ...
    # Arc to p2 via p1
    def threePointArc(self, p1: Tuple[float, float], p2: Tuple[float,
                                                                float]) -> Self: ...

    # Spline through given points
    def splineTo(self, *p: Tuple[float, float]) -> Self: ...
    # Close loop (make a line back to start point)
    def close(self) -> Self: ...
    # Draw circle with center at current point
    def circle(self, radius: float) -> Self: ...
    # assign tag to the current point
    def pointTag(self, tag: str) -> Self: ...
    # assign tag to the current curve
    def curveTag(self, tag: str) -> Self: ...

class Profile:
    # Create a face with a optional tag for reference
    def __init__(self, tag: Optional[str] = None) -> None: ...
    # Add one or more loops to the face
    def addLoop(self, *loops: Loop) -> None: ...

class Sketch:
    # plane format: {"origin": [x, y, z], "x_axis": [x, y, z], "normal": [x, y, z]}

    def __init__(self, plane: Dict) -> None: ...
    def addProfile(self, *profiles: Profile) -> None: ...

class CADShape(ABC):
    def __init__(self) -> None: ...
    # entities: list of dicts, each with {"capType": "START"|"END"|"
                                         "SWEEP", 'referenceId': str},
                                         specifying referenced entities

    def Chamfer(self, width: float, entities: List[Dict]) -> Self: ...
    def Fillet(self, radius: float, entities: List[Dict]) -> Self: ...
    def Shell(self, thickness: float, entities: List[Dict]) -> Self: ...
    ...
    def union(self, shape: CADShape) -> Self: ...
    def cut(self, shape: CADShape) -> Self: ...
    def intersect(self, shape: CADShape) -> Self: ...

class Extrude(CADShape):
    def __init__(self, sketch: Sketch, distance: Union[float, Tuple[
                                                                float, float]]) -> None: ...

class Revolve(CADShape):
```

```
# axis format: {"point": [x, y, z], "direction": [x, y, z]}, axis
#               is defined by a point and a
#               direction
def __init__(self, sketch: Sketch, axis: Dict, angle: Union[float,
    Tuple[float, float]]) -> None:
    ...
```

A.4 RAG Settings

To equip the SSR-based CAD corpus locally for the retrieval-augmented generation of Seek-CAD, we use Dify [6] in 0.15.3 version, which is a docker-based [7] platform. For each query, we use bge-m3 [5] to embed it into a vector that is used to calculate similarity score with CAD models stored in the local SSR-based corpus. Practically, we select Top-3 samples based on a hybrid searching strategy (Sec 3.1) for the augmentation retrieval. To incorporate the retrieved content into the input for Seek-CAD, we apply a chunk-based prompt construction strategy. Each retrieved pair $(d_j, s_j) \in \mathcal{R}_T$ is formatted as:

$$[\text{Chunk}_j] := \text{"Description: " } d_j \parallel \text{"CAD code: " } s_j, \quad (11)$$

where $\oplus$ denotes string concatenation. The final prompt for Seek-CAD becomes:

$$T \oplus (d_j, c_j) = \text{"Query: " } T \oplus \text{Chunk}_1 \oplus \dots \oplus \text{Chunk}_k. \quad (12)$$

A.5 Novel Visualizations

Benefiting from the prior knowledge of DeepSeek-R1, Seek-CAD is capable of generating not only industry-oriented CAD models but also non-industrial ones that differ in style from those in our local CAD corpus. As shown in Figure 9, this demonstrates that Seek-CAD does not rely entirely on our local CAD corpus and is able to generate models with a certain degree of novelty.




Figure 9: Novel showcases generated by Seek-CAD, which are different from the style of CAD models in the local CAD corpus.

A.6 Enlarged Version of Showcases in Figure 5(c) of the Main Manuscript

We provide an enlarged version of showcases in Figure 5(c) for a clear view as shown in Figure 10.




Figure 10: The enlarged version of showcases in Figure 5(c) of the main manuscript.

A.7 More showcases of CAD Editing




Figure 11: The showcases of similar generations.

A.8 More Showcases of Similar Generations





Figure 12: The showcases of similar generations.

A.9 The Detail Showcase of Complete Generation Process

Here we provide two demo showcases of complete generation process by Seek-CAD, including the input query T , CoT from DeepSeek-R1, initial generated CAD models, feedback content, and finalized CAD model (Figure 13).

(1) Demo



(2) Demo




Figure 13: The showcases of complete generation process.

A.10 More showcases by Seek-CAD




Figure 14: The more generated showcases by Seek-CAD.