

Revisiting CAD Model Generation by Learning Raster Sketch

Pu Li^{1,2*}, Wenhao Zhang^{1*}, Jianwei Guo^{3,1}, Jinglu Chen¹, Dong-Ming Yan^{1,2†}

¹MAIS, Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³School of Artificial Intelligence, Beijing Normal University

{lipu2021, wenhao.zhang, chenjinglu}@ia.ac.cn, jianwei.guo@bnu.edu.cn, yandongming@gmail.com

Abstract

The integration of deep generative networks into generating Computer-Aided Design (CAD) models has garnered increasing attention over recent years. Traditional methods often rely on discrete sequences of parametric line/curve segments to represent sketches. Differently, we introduce RECAD, a novel framework that generates **R**aster sketches and **3D Extrusions** for CAD models. Representing sketches as raster images offers several advantages over discrete sequences: 1) it breaks the limitations on the types and numbers of lines/curves, providing enhanced geometric representation capabilities; 2) it enables interpolation within a continuous latent space; and 3) it allows for more intuitive user control over the output. Technically, RECAD employs two diffusion networks: the first network generates extrusion boxes conditioned on the number and types of extrusions, while the second network produces sketch images conditioned on these extrusion boxes. By combining these two networks, RECAD effectively generates sketch-and-extrude CAD models, offering a more robust and intuitive approach to CAD model generation. Experimental results indicate that RECAD achieves strong performance in unconditional generation, while also demonstrating effectiveness in conditional generation and output editing.

Introduction

The digital genesis of modern artifacts, from everyday consumer products to complex industrial machinery, is now deeply intertwined with Computer-Aided Design (CAD) systems. Central to many CAD workflows is sketch-based modeling, where 2D sketches imbued with geometric constraints and design intent are transformed into intricate 3D models through a series of feature-based modeling operations, ultimately giving rise to complex assemblies.

Among these feature-based modeling operations, extrusion is the most prevalent, allowing designers to generate 3D shapes by extending 2D sketches along a defined path. However, the traditional parametric approach to sketch extrusion, while powerful, presents significant challenges. Creating these parametric models, where 2D sketches are extruded into 3D forms based on numerous parameters and

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

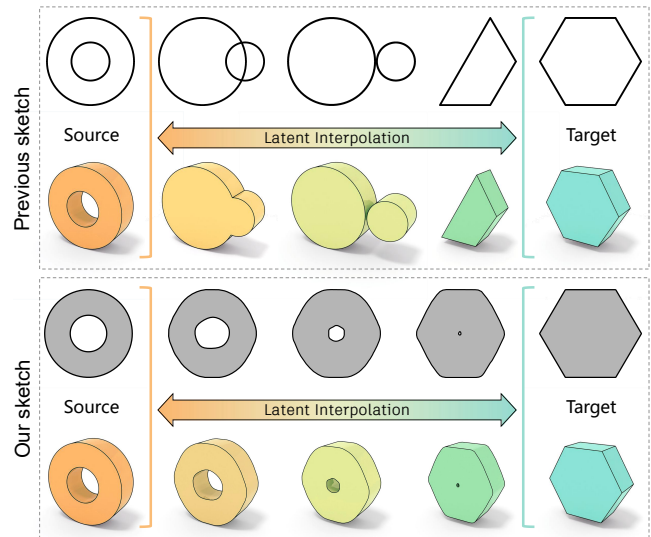


Figure 1: Top: Previous approaches (Wu, Xiao, and Zheng 2021; Xu et al. 2022, 2023; Wang et al. 2024) represent sketches using vector sequences, which often fail to produce smooth transitions in shape and topology during latent space interpolation. Bottom: RECAD learns from rasterized sketches, ensuring more plausible interpolated shapes, which results in more natural and continuous transformations between different sketches.

equations, is tedious and expert-intensive. Furthermore, the rigid nature of these parametric dependencies makes major design modifications difficult, often causing models to fail. This lack of flexibility hinders iterative design and highlights the need for more intelligent CAD systems.

With the advance of deep neural networks, learning-based methods have emerged as a promising avenue for generating diverse CAD models. A prevalent approach involves training generative models by leveraging construction sequences as supervision (Wu, Xiao, and Zheng 2021; Xu et al. 2022, 2023; Wang et al. 2024). These methods typically represent sketch-and-extrude sequences as parameterized discrete sequences, enabling the application of powerful sequence learning techniques to the task of CAD model generation.

Despite the demonstrated success of sequence-based




Figure 2: Sample data from DeepCAD dataset. Previous methods, constrained by sequence length limitations, were restricted to learning a subset of simpler shapes within the dataset (gray). Our raster-based sketch representation enables the learning of more complex shapes (green).

learning methods for CAD model generation, several limitations remain: 1) The reliance on sequences introduces a bottleneck, limiting the complexity of representable sketches and the number of commands that can be effectively utilized. This constraint restricts the generation of intricate models. 2) Current methods often generate semantically invalid command sequences for CAD modeling, resulting in non-viable models and hindering the reliability and usability of the output. 3) Representing sketches as discrete parameter sequences makes generating smooth and continuous shape transformations difficult, leading to unnatural results for latent space interpolation, as illustrated in Fig.1. This poses a challenge for generating plausible sketches, for which plausible interpolation is essential (Radford, Metz, and Chintala 2015; Goodfellow et al. 2014; Higgins et al. 2017).

In this paper, we revisit the CAD generation task and propose a novel representation that follows the sketch-and-extrude paradigm without using command sequences. Instead of representing sketches as curve sequences, we leverage raster images. Extrusion parameters, such as origin, direction, and Boolean operation type, are implicitly determined by the dimensions and type of the extrusion box. This new representation addresses the limitations of sequence-based methods in several key ways. Using raster images to represent sketches allows for depicting highly complex shapes, overcoming the bottleneck imposed by discrete parameter sequences (e.g., fixed maximum length of curves per loop), as shown in Fig.2. Additionally, the intuitive nature of image-based sketches reduces the likelihood of generating semantically or geometrically invalid commands, leading to more reliable and usable outputs. Moreover, representing sketches as images facilitates the generation of smooth and continuous shape transformations, unlike the discrete nature of previous methods. Finally, by employing extrusion boxes, our method provides a concise and intuitive way to control the topological relationships between extruded shapes, further enhancing the modeling process.

Building upon this novel sketch-extrusion representation, we present RECAD, a novel generative framework that generates sketch-extrusion pairs using Denoising Diffusion Probabilistic Models (DDPM) (Ho, Jain, and Abbeel 2020). Concretely, RECAD comprises a sketch image VAE and two diffusion networks. A Variational Autoencoder (VAE) (Kingma and Welling 2013) is utilized to learn the latent features of raster sketches. Two diffusion networks are em-

ployed in a sequential generation process. The first diffusion network generates the coordinates defining the extrusion boxes. Conditioned on these boxes, the second network generates latent features representing the raster sketches. By extruding sketches into 3D solids and integrating the topology defined by the extrusion boxes, we achieve the final CAD model.

Leveraging the adaptable architecture of RECAD, our work explores advancements in controllable generation. Beyond the typical autocompletion capabilities, we introduce a novel feature: extrusion box conditional generation. This allows for automatically generating sketches based on a user-defined structure of extrusion boxes, resulting in a detailed 3D model. Additionally, RECAD allows users to easily edit generated sketches using simple image editing tools, such as monochrome brushes, to modify local geometry in a natural and intuitive manner. In summary, our contributions include:

- A novel sketch-extrusion representation leverages raster sketch and extrusion box pairs for CAD model construction.
- A latent diffusion generative framework that enables random generation of complicated CAD models, auto-completion from partial input, extrusion box conditional generation, and flexible user editing.
- Extensive experiments demonstrate the superiority of RECAD in CAD generation through comprehensive comparisons with existing state-of-the-art approaches.

Related Work

Constructive solid geometry. Constructive Solid Geometry (CSG) represents complex 3D shapes through Boolean operations on simple geometric primitives. This compact representation has been employed for CAD reconstruction tasks with program synthesis (Du et al. 2018; Nandi et al. 2017, 2018). For instance, CSGNet (Sharma et al. 2018) uses reinforcement learning to minimize reconstruction errors. UCSGNet (Kania, Zieba, and Kajdanowicz 2020) takes this a step further by eliminating the need for ground truth CSG trees during inference. Other works also explored unsupervised learning with specialized parametric primitives (Chen, Tagliasacchi, and Zhang 2020; Ren et al. 2021; Yu et al. 2022, 2023). While deep learning methods hold promise for advancing CSG-based 3D model generation, they are still in the exploration phase, with limited datasets posing a significant challenge.

B-rep CAD generation. Recent research endeavors have sought to directly generate Boundary Representation (B-rep) CAD models. Direct B-rep generation involves constructing the necessary geometric and topological elements, trimming surface patches, and joining them into a unified solid model. Leveraging learning-based methods, previous works have made advances in generating parametric curves (Wang et al. 2020) and surfaces (Sharma et al. 2020; Li et al. 2023b). To replace fixed topological templates (Smirnov, Bessmeltsev, and Solomon 2019), creating generalized topology for sketches (Willis et al. 2021) and solid models (Wang, Zheng, and Zhou 2022; Guo et al. 2022; Jayaraman et al. 2022; Xu et al. 2024) is then explored.

Sketch-and-extrude CAD generation. Sequential CAD modeling operations combined with learning-based methods have made significant progress on several tasks. Some works progress in generating 2D engineering sketches (Willis et al. 2021; Para et al. 2021; Ganin et al. 2021; Seff et al. 2021) in editable parametric CAD files obtained from the generated sequences, and converting between command sequences and hand drawings (Li et al. 2020, 2022; Hähnlein et al. 2022). DeepCAD (Wu, Xiao, and Zheng 2021) produces a large-scale dataset of CAD models in the form of sketch-and-extrude sequences and first learns to generate editable CAD models. Later works further improve the performance of shape generation by employing auto-regressive generative models (Xu et al. 2022), designing a hierarchical tree of neural codes mastering both high level concepts and local level geometry (Xu et al. 2023), and integrating discrete diffusion models (Wang et al. 2024). In the field of 3D model reconstruction, approaches learning 2D sketches and predicting the corresponding extrusions to reconstruct the CAD models in supervised (Uy et al. 2022), and unsupervised (Li et al. 2023a; Ren et al. 2022; Li et al. 2024) manners are proposed.

This work explores CAD model generation, focusing specifically on the sketch-and-extrude modeling paradigm. Departing from prior art that relies on discrete parametric representations of sketches, we propose leveraging raster images as a more flexible and expressive alternative. Furthermore, we encapsulate all extrusion information within an extrusion box.

Generative diffusion models. DDPM (Ho, Jain, and Abbeel 2020) are good at generative tasks, especially in image generation and restoration (Ho, Jain, and Abbeel 2020; Saharia et al. 2022; Lugmayr et al. 2022). Latent diffusion models (LDM) operate in a low-dimensional latent space rather than original data space, resulting in more efficient sampling and promising generative performance on numerous tasks such as high-resolution image synthesis (Vahdat, Kreis, and Kautz 2021; Rombach et al. 2022). Diffusion models also show the ability to generate geometry with (Shabani, Hosseini, and Furukawa 2023; Liu et al. 2024) or without topology as input conditioning (Alliegro et al. 2023; Xu et al. 2024). Borrowing architecture from VQ-Diffusion (Gu et al. 2022) which combines VQ-VAE (Van Den Oord, Vinyals et al. 2017) and conditional DDPM, VQ-CAD (Wang et al. 2024) integrates discrete diffusion models for CAD model generation. Our work includes two diffusion networks for generating 3D extrusion boxes and raster sketches representing detailed 2D geometry.

Representation and Overview

In this section, we first introduce our novel representation using raster sketches paired with extrusion boxes. We then provide an overview of our proposed approach for generating CAD models based on this representation.

Novel Sketch-Extrusion Representation

Sequence-based representations, while commonly used in CAD model generation, limit design complexity, semantic validity, and smooth shape transformations. To address these

challenges, we propose a novel sketch-extrusion paradigm that leverages raster images for sketch representation, allowing for greater complexity and smoother transformations. Additionally, we introduce extrusion boxes, establishing a direct mapping between the 2D sketch and the 3D extrusion. This representation simplifies the extrusion parameters themselves into a few key controls within the extrusion box, allowing for intuitive manipulation of the resultant topology. Below, we detail our raster sketch representation and the extrusion box concept.

Raster sketch. A raster sketch is a binary image depicting the shape of the sketch. Pixel values of 1 correspond to areas suitable for extrusion, while 0 represents empty space. This image-based representation allows for the portrayal of complex geometries without concern for the count of internal loops and curves. Moreover, compared to the sequence-based representation, this pictorial format facilitates more straightforward and valid interpolation within the latent space.

Extrusion box. The DeepCAD dataset defines extrusion using nine continuous and two discrete parameters. We note that the extrusion direction for most models aligns with the principal axes. Furthermore, the symmetry of extrusion allows us to consolidate bidirectional extrusion into a unidirectional representation. Motivated by these observations, we propose a simplified and more straightforward representation called the “extrusion box”. An extrusion box is fully defined by eight parameters:

$$E = (x_{min}, y_{min}, z_{min}, x_{max}, y_{max}, z_{max}, \omega, \mathcal{B}). \quad (1)$$

The first three parameters represent the bottom-left corner of the extrusion box, while the subsequent three indicate the top-right corner. The term ω specifies the extrusion direction along one of the principal axes (i.e., x, y, or z). \mathcal{B} denotes one of the three Boolean operations (i.e., union, subtraction intersection, or intersection). This representation directly couples the 2D sketch to the 3D extrusion: The dimensions of the sketch are determined by the ranges of the two axes perpendicular to the extrusion direction, while the range along the extrusion direction dictates the extrusion height.

Each pair of raster sketch and extrusion box can naturally form a solid through the sketch-extrude process. Multiple sketch-extrusion pairs can link the generated solids topologically to express complex CAD models.

Overview

We aim to generate sketch-extrusion pairs in the format of our novel representation and convert them into parametric CAD models. In RECAD, a sequential two-stage denoising is utilized for shape generation, as illustrated in Fig.3. In the first stage, a set of extrusion topology guidance E_g comprising one-hot encoded direction and operation is randomly sampled from its ground-truth distribution. We then use a Transformer-based denoiser to denoise a random Gaussian noise conditioned on E_g , and obtain the extrusion coordinates E_c . A set of complete extrusion boxes E are formed by combining E_g and E_c . In the second stage, a UNet-based denoiser is employed to generate latent sketches S_z




Figure 3: Framework of RECAD. A sketch image VAE, shown on the left, learns a latent feature base of rasterized input sketches. This feature base is then utilized for efficient sketch retrieval. The right panel details the two-stage diffusion process. First, extrusion boxes are progressively generated conditioned on user-specified operation and direction. Subsequently, sketch features are generated conditioned on these bounding boxes and then decoded into raster sketches using the VAE decoder. $E^{(t)}$ and $S^{(t)}$ represent visualizations of the extrusion boxes and the sketch-extrusion result at step t , respectively.

conditioned on the complete extrusion boxes E . The latent sketches S_z are then decoded to sketch images S through a sketch image VAE. We finally combine the generated extrusion boxes and sketch images to construct CAD models.

Method

Here, we present RECAD, a novel generative framework comprising a sketch image VAE and two diffusion networks.

Sketch Image VAE

We employ a sketch image VAE to compress sketch images into latent embeddings as in Stable Diffusion (Rombach et al. 2022). The backbone of the sketch image VAE is a 2D convolution UNet (Ronneberger, Fischer, and Brox 2015). The grayscale images to be used as VAE input are all resized to a size of 32×32 . Concretely, given a binary sketch image $S \in \mathbb{R}^{32 \times 32 \times 1}$, the VAE encoder \mathcal{E} encodes S into a latent representation by a downsampling factor of 8 and a feature depth of 3, yielding $S_z \in \mathbb{R}^{4 \times 4 \times 3}$. The VAE decoder \mathcal{D} then reconstructs the sketch image $\tilde{S} = \mathcal{D}(S_z) = \mathcal{D}(\mathcal{E}(S))$. The training process of the sketch image VAE involves optimizing two components: the mean squared error (MSE) for reconstruction loss between S and \tilde{S} , and the Kullback-Leibler (KL) divergence as a regularization term.

Two-Stage Diffusion

Diffusion models are generative models that learn to gradually denoise a sample from Gaussian noise towards a data distribution, consisting of a forward diffusing process that adds noise and a reverse denoising process that adds structure. Our two-stage diffusion contains a Transformer-based extrusion denoiser and a UNet-based sketch denoiser. We train the denoisers with topologically organized extrusion boxes and latent sketches with inherent geometry details, following the training strategy in DDPM (Ho, Jain, and Abbeel 2020).

Forward diffusing process. During training, the forward process takes a data sample x_0 , which could be extrusion coordinates or latent sketch features in our case, and gradually adds Gaussian noise to the data in T time steps. Each forward step is computed by:

$$q(x_t|x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}\right), \quad (2)$$

where β_t is a noise variance schedule. At a specific time step t , the noisy sample features x_t are sampled as

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t \quad (3)$$

using the notation $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, $\alpha_t = 1 - \beta_t$ and Gaussian noise $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$. After T time steps, the data sample will be completely corrupted into a random noise. The latent sketches will lose their inherent shape information, and the coordinate vectors of extrusions will be scrambled, resulting in the forgetting of relative topology.

Two-stage denoising process. We then train DDPM to reverse the diffusing process. The reverse process denoises a pure Gaussian noise step by step until recovering x_0 and learns neural networks to directly predict the noise-removed data distribution of extrusion boxes and latent sketches.

We apply a sequential two-stage denoising, which first generates extrusion boxes conditioned on ground-truth extrusion guidance and then generates latent sketches conditioned on these extrusion boxes.

At the beginning of the first stage, we randomly sample a set of extrusion guidance $E_g \in \mathbb{R}^5$ from its ground-truth distribution E_g^* , which is derived from a histogram of extrusion directions and Boolean operations in our training data:

$$E_g = (\Phi(\omega, 3), \Phi(\mathcal{B}, 2)) \sim P(E_g^*), \quad (4)$$

where Φ denotes one-hot encoding. Boolean operation \mathcal{B} is encoded as a 2D vector, encompassing only union and difference. The vector E_g acts as the conditioning input for the denoiser with a Transformer backbone (Vaswani 2017), which

processes the noisy extrusion box coordinates $E_c \in \mathbb{R}^6$. E_g and E_c are embedded with the same feature dimension:

$$\mathbf{E}_g = \text{MLP}(W_g E_g), \mathbf{E}_c = \text{MLP}(W_c E_c), \quad (5)$$

where $W_g \in \mathbb{R}^{n \times 5}$ and $W_c \in \mathbb{R}^{n \times 6}$ are embedding matrices of n -dimension. MLP is fully connected layers with SiLU activation. We employ direct condition injection in this stage: $\tilde{\mathbf{E}}_c \leftarrow \mathbf{E}_c + \mathbf{E}_g$. The first stage produces a set of denoised extrusion coordinates, $E_c^{de} \in \mathbb{R}^6$. These coordinates are concatenated with the corresponding extrusion guidance, E_g , forming the complete extrusion boxes, $E \in \mathbb{R}^{11}$:

$$E = E_c^{de} \parallel E_g. \quad (6)$$

E serves as the condition for the second denoiser.

In the second stage, inspired by the exceptional performance of video diffusion models (Ho et al. 2022) in generating sequential images, we similarly employ a 3D conditional UNet (Çiçek et al. 2016) as the backbone of the denoiser for sketch sequence generation. Each extrusion box E is embedded as $\mathbf{E} = \text{MLP}(WE)$ where $W \in \mathbb{R}^{d \times 11}$ likewise as in the first stage. The input to the second denoiser is latent sketch features $S_z \in \mathbb{R}^{4 \times 4 \times 3}$. In this stage, we apply cross-attention to control the denoising UNet with \mathbf{E} as a conditioning feature. A set of sketch-extrusion pairs are obtained at the end of the two-stage denoising.

Note that time embedding is also passed to the denoiser in each stage. Overall, the distribution of generated extrusion-sketch pairs can be defined as the product of conditional distributions $p(E_c|E_g)$ and $p(S_z|E)$.

Loss function. Following DDPM (Ho, Jain, and Abbeel 2020), we train our Transformer-based denoiser and 3D UNet-based denoiser respectively with the same simple L2-norm regression loss,

$$L = \mathbb{E}_{t, x_0, \epsilon_t} \left[\left\| \epsilon_t - \epsilon_\theta \left(\sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t, t \right) \right\|^2 \right], \quad (7)$$

where $\epsilon_\theta \left(\sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t, t \right)$ can be simplified to $\epsilon_\theta(x_t, t)$ using Eq. 3. ϵ_t is the ground-truth Gaussian noise at time step t from the forward diffusing process and is compared against $\epsilon_\theta(x_t, t)$. During sampling, the inferred noise at each time step is utilized to progressively denoise a random Gaussian noise, refining it towards a cleaner signal.

Extrusion with Raster Sketch

The conversion from raster sketches to contours is essential for extruding them into 3D solids. We present two methods for this conversion:

Contour extraction. We utilize the Teh-Chin chain approximation (Teh and Chin 1989) to convert binary raster images into polygonal contours. To streamline the CAD generation process, we simplify these contours using the Douglas-Peucker algorithm (Douglas and Peucker 1973). The simplified contours are then extruded to 3D prisms.

Sketch retrieval. Our sketch image VAE embeds each training sketch into a feature space. During testing, we retrieve the closest matching sketch from the database based

on the Euclidean distance between the generated feature vector of the sketch and the database entries. This retrieved sketch provides predefined contour curves, which are then scaled and extruded to form the 3D solid. Compared to directly extracting contours, this method may limit the diversity of the sketch contours. However, it offers cleaner and more coherent curves, bypassing the process of contour extraction and simplification, thus improving efficiency. We employ sketch retrieval for all experiments except those involving user edits.

Implementation Details

Models are implemented in PyTorch and trained on four NVIDIA RTX A6000 GPUs. We use an AdamW (Loshchilov and Hutter 2017) optimizer with the learning rate $5e-4$ for optimization. The sketch Image VAE is trained for 500 epochs at batch size 512 and two denoisers are trained for 2000 epochs at batch size 256. During inference, PNDM (Liu et al. 2022) is employed for fast sampling. More details are provided in the supplementary material.

Experiments

In this section, we demonstrate the effectiveness of our method on two kinds of generation tasks: 1) unconditional generation and 2) controllable generation including auto-completion from partial input, extrusion box conditional generation, and image-based sketch editing.

Setup

Dataset preparation. We utilize the extensive DeepCAD dataset (Wu, Xiao, and Zheng 2021) which contains CAD models with detailed sequences of sketch-and-extrude modeling operations. We follow the procedures of prior works (Willis et al. 2021; Xu et al. 2022, 2023) to detect and eliminate duplicate models within the training set.

Evaluation metrics. Three metrics evaluate the discrepancy between generated models and the reference test set (Wu, Xiao, and Zheng 2021):

- *Coverage (COV)* measures the diversity of generated shapes by calculating the percentage of ground-truth models that are matched by at least one generated shape using Chamfer Distance (CD).
- *Minimum Matching Distance (MMD)* evaluates the fidelity of generated shapes by averaging the CD between each ground-truth model and its closest generated match.
- *Jensen-Shannon Divergence (JSD)* reports the similarity between the distributions of the reference test set and generated set.

To measure the robustness and uniqueness of generated results, the following metrics are computed (Xu et al. 2024):

- *Valid* is the percentage of raw model outputs successfully converted into valid Boundary Representations.
- *Novel* is the percentage of generated data that do not exist in the training set.
- *Unique* is the percentage of generated data found only once in the generated set.




Figure 4: Visual comparison on the DeepCAD dataset. Each row, from top to bottom, shows shapes generated with 1 to 5 extrusion operations, respectively. Different colors are used for clear visual distinction.

Method	COV ↑(%)	MMD ↓	JSD ↓	Valid ↑(%)	Novel ↑(%)	Unique ↑(%)
DeepCAD	81.12	0.921	0.995	86.24	91.7	85.8
SkexGen	83.39	0.808	0.632	86.73	99.1	99.8
HNC-CAD	85.33	0.772	0.558	76.07	93.9	99.7
Ours (CE)	86.67	0.764	0.616	95.88	94.5	98.6
Ours	86.67	0.768	0.593	99.65	94.4	98.3

Table 1: Quantitative comparison of unconditional CAD generation. The second-to-last row reports the results of RECAD with contour extraction. Both MMD and JSD are multiplied by 10^2 .

Unconditional CAD Generation

In this section, we evaluate our models on unconditional CAD generation task. Sketch-extrusion pairs are generated by our diffusion models from random Gaussian noise and then converted to CAD models. We compare our generated samples with state-of-the-art methods, including DeepCAD (Wu, Xiao, and Zheng 2021), SkexGen (Xu et al. 2022), and HNC-CAD (Xu et al. 2023). We generate 10,000 CAD models using each method and evaluate them against 2,500 ground-truth models randomly sampled from the reference test set in each run. We repeated this evaluation process three times and averaged the metric values across these runs to obtain the final results.

Quantitative evaluation. As reported in Tab. 1, RECAD outperforms other methods in terms of COV and MMD. This suggests that RECAD is capable of producing more diverse parametric CAD models. While our method achieves a slightly worse JSD score compared to HNC-CAD, it still exceeds SkexGen. RECAD also achieves competitive Novel and Unique scores. Furthermore, our method significantly surpasses all baseline models in terms of Valid, highlighting the robustness of our image-based sketch representation and the effectiveness of our generative framework.




Figure 5: Autocompleted CAD models (green) from partial inputs (gray).

Qualitative evaluation. Fig. 4 presents qualitative comparisons among different methods across varying numbers of extrusion operations. A key advantage of our image-based representation lies in its ability to capture the complexity of input sketches, which is particularly evident when dealing with a limited number of extrusions. With one or two extrusion operations, other methods produce relatively simple models, while RECAD generates more intricate designs. As the number of extrusions increases, RECAD excels in generating complex sketches while ensuring logical coherence between different extruded parts, maintaining the authenticity of the CAD model. We include additional results and shape novelty analysis that compares generated models with their closest training samples in the supplementary material.

Controllable CAD Generation

We demonstrate controllable generation results of RECAD in three application scenarios for CAD design.

Autocompletion from partial input. A set of sketch-extrusion pairs with an amount below the extrusion box limit is considered a partial input. Fig. 5 shows that our pre-trained model can automatically generate the missing parts based on existing extrusion boxes and sketches, resulting in more complex CAD models. The autocompletion process involves two stages: extrusion box completion and sketch completion. Drawing inspiration from ReShape (Lugmayr et al. 2022) and BrepGen (Xu et al. 2024), at each time step




Figure 6: Extrusion box conditional generation. Different colors are used for clear visual distinction.

of the denoising process, we simultaneously diffuse the partial input to the corresponding time step and replace a randomly selected portion of the tokens being denoised with the diffused partial input. The completed extrusion boxes obtained in the first stage are used as conditions to assist in generating sketches in the second stage.

Extrusion box conditional generation. We observe that the extrusion boxes significantly influence the overall structure of the CAD model, sometimes even suggesting its intended function. Inspired by layout-conditioned image and 3D scene generation (Zheng et al. 2023; Bahmani et al. 2023), we explore CAD generation conditioned on the extrusion boxes. As shown in Fig. 6, we manually create extrusion boxes (see the first column) and then use RECAD to generate sketches, resulting in various plausible outcomes. Interestingly, although trained primarily on mechanical parts from the DeepCAD dataset, our model generated reasonable shapes even when presented with extrusion boxes representing common objects like chairs and tables. This approach streamlines CAD creation and grants designers efficient, high-level control over the generated model structure.

Image-based sketch editing. A key benefit of employing binary raster sketches in RECAD is the intuitiveness of the editing process. As demonstrated in Fig. 7, users can select the extrusion target and modify the sketch directly with simple black or white brush strokes within any standard image editor. This eliminates the need for specialized CAD software, making 3D CAD model creation and modification remarkably user-friendly.

Limitations and Future Work

While image-based sketches in RECAD provide a natural interface for design, their dependence on raster images can impede the accurate representation of intricate contours, particularly in low-resolution contexts. Moreover, current CAD model generation methods, including RECAD, lack




Figure 7: Examples of user edits to the generated results. After selecting the target extrusion, users can directly modify the appearance of the 3D model through simple doodling.

built-in mechanisms to ensure the plausibility and realism of the generated designs. Introducing adversarial losses or human-in-the-loop feedback mechanisms could be a promising direction for future research. Additionally, given the widespread use of image-based representations over curve sequences, exploring the integration of RECAD with complementary modalities such as text descriptions or semantic segmentation could further enhance the design process.

Conclusion

We have presented a novel generative diffusion framework for CAD model generation. As the key to our approach, we have introduced a novel sketch-extrusion representation consisting of raster sketch and extrusion box pairs to substitute the predominant sequence-based representation. A two-stage denoising process is adopted to sequentially generate sketch-extrusion pairs which are then converted to CAD models. Extensive quantitative and qualitative evaluations show that RECAD surpasses previous approaches, generating CAD models with greater diversity and complexity. Furthermore, RECAD offers users an intuitive and visual interface for modifying both the global topology and local geometry of CAD models, making the editing process more accessible and aligned with user expectations.

Acknowledgments

This work is partially funded by the National Natural Science Foundation of China (12494550, 12494553, 62172416 and 62172415), the Strategic Priority Research Program of the Chinese Academy of Sciences (XDB0640000 and XDB0640200), the Beijing Natural Science Foundation (Z240002), and the Guangdong Basic and Applied Basic Research Foundation (2023B1515120026).

References

- Alliegro, A.; Siddiqui, Y.; Tommasi, T.; and Nießner, M. 2023. Polydiff: Generating 3d polygonal meshes with diffusion models. *arXiv preprint arXiv:2312.11417*.
- Bahmani, S.; Park, J. J.; Paschalidou, D.; Yan, X.; Wetstein, G.; Guibas, L.; and Tagliasacchi, A. 2023. Cc3d: Layout-conditioned generation of compositional 3d scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7171–7181.
- Chen, Z.; Tagliasacchi, A.; and Zhang, H. 2020. Bsp-net: Generating compact meshes via binary space partitioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 45–54.
- Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S. S.; Brox, T.; and Ronneberger, O. 2016. 3D U-Net: learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*, 424–432. Springer.
- Douglas, D. H.; and Peucker, T. K. 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: the international journal for geographic information and geovisualization*, 10(2): 112–122.
- Du, T.; Inala, J. P.; Pu, Y.; Spielberg, A.; Schulz, A.; Rus, D.; Solar-Lezama, A.; and Matusik, W. 2018. Inversecsg: Automatic conversion of 3d models to csg trees. *ACM Transactions on Graphics (TOG)*, 37(6): 1–16.
- Ganin, Y.; Bartunov, S.; Li, Y.; Keller, E.; and Saliceti, S. 2021. Computer-aided design as language. *Advances in Neural Information Processing Systems*, 34: 5885–5897.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gu, S.; Chen, D.; Bao, J.; Wen, F.; Zhang, B.; Chen, D.; Yuan, L.; and Guo, B. 2022. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10696–10706.
- Guo, H.; Liu, S.; Pan, H.; Liu, Y.; Tong, X.; and Guo, B. 2022. Complexgen: Cad reconstruction by b-rep chain complex generation. *ACM Transactions on Graphics (TOG)*, 41(4): 1–18.
- Hähnlein, F.; Li, C.; Mitra, N. J.; and Bousseau, A. 2022. Cad2sketch: Generating concept sketches from cad sequences. *ACM Transactions on Graphics (TOG)*, 41(6): 1–18.
- Higgins, I.; Matthey, L.; Pal, A.; Burgess, C. P.; Glorot, X.; Botvinick, M. M.; Mohamed, S.; and Lerchner, A. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR (Poster)*, 3.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video diffusion models. *Advances in Neural Information Processing Systems*, 35: 8633–8646.
- Jayaraman, P. K.; Lambourne, J. G.; Desai, N.; Willis, K. D.; Sanghi, A.; and Morris, N. J. 2022. Solidgen: An autoregressive model for direct b-rep synthesis. *arXiv preprint arXiv:2203.13944*.
- Kania, K.; Zieba, M.; and Kajdanowicz, T. 2020. UCSG-NET-unsupervised discovering of constructive solid geometry tree. *Advances in neural information processing systems*, 33: 8776–8786.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Li, C.; Pan, H.; Bousseau, A.; and Mitra, N. J. 2020. Sketch2cad: Sequential cad modeling by sketching in context. *ACM Transactions on Graphics (TOG)*, 39(6): 1–14.
- Li, C.; Pan, H.; Bousseau, A.; and Mitra, N. J. 2022. Free2cad: Parsing freehand drawings into cad commands. *ACM Transactions on Graphics (TOG)*, 41(4): 1–16.
- Li, P.; Guo, J.; Li, H.; Benes, B.; and Yan, D.-M. 2024. SfmCAD: Unsupervised CAD Reconstruction by Learning Sketch-based Feature Modeling Operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4671–4680.
- Li, P.; Guo, J.; Zhang, X.; and Yan, D.-M. 2023a. Secadnet: Self-supervised cad reconstruction by learning sketch-extrude operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16816–16826.
- Li, Y.; Liu, S.; Yang, X.; Guo, J.; Guo, J.; and Guo, Y. 2023b. Surface and edge detection for primitive fitting of point clouds. In *ACM SIGGRAPH 2023 conference proceedings*, 1–10.
- Liu, J.; Tam, H. I. I.; Mahdavi-Amiri, A.; and Savva, M. 2024. CAGE: Controllable Articulation GEneration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17880–17889.
- Liu, L.; Ren, Y.; Lin, Z.; and Zhao, Z. 2022. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11461–11471.
- Nandi, C.; Caspi, A.; Grossman, D.; and Tatlock, Z. 2017. Programming language tools and techniques for 3D printing. In *2nd Summit on Advances in Programming Languages (SNAPL 2017)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
- Nandi, C.; Wilcox, J. R.; Panchekha, P.; Blau, T.; Grossman, D.; and Tatlock, Z. 2018. Functional programming for compiling and decompiling computer-aided design. *Proceedings of the ACM on Programming Languages*, 2(ICFP): 1–31.

- Para, W.; Bhat, S.; Guerrero, P.; Kelly, T.; Mitra, N.; Guibas, L. J.; and Wonka, P. 2021. Sketchgen: Generating constrained cad sketches. *Advances in Neural Information Processing Systems*, 34: 5077–5088.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Ren, D.; Zheng, J.; Cai, J.; Li, J.; Jiang, H.; Cai, Z.; Zhang, J.; Pan, L.; Zhang, M.; Zhao, H.; et al. 2021. Csg-stump: A learning friendly csg-like representation for interpretable shape parsing. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12478–12487.
- Ren, D.; Zheng, J.; Cai, J.; Li, J.; and Zhang, J. 2022. Extrudenet: Unsupervised inverse sketch-and-extrude for shape parsing. In *European Conference on Computer Vision*, 482–498. Springer.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4): 4713–4726.
- Seff, A.; Zhou, W.; Richardson, N.; and Adams, R. P. 2021. Vitruvion: A generative model of parametric cad sketches. *arXiv preprint arXiv:2109.14124*.
- Shabani, M. A.; Hosseini, S.; and Furukawa, Y. 2023. Housediffusion: Vector floorplan generation via a diffusion model with discrete and continuous denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5466–5475.
- Sharma, G.; Goyal, R.; Liu, D.; Kalogerakis, E.; and Maji, S. 2018. Csgnet: Neural shape parser for constructive solid geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5515–5523.
- Sharma, G.; Liu, D.; Maji, S.; Kalogerakis, E.; Chaudhuri, S.; and Měch, R. 2020. Parsenet: A parametric surface fitting network for 3d point clouds. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, 261–276. Springer.
- Smirnov, D.; Bessmeltsev, M.; and Solomon, J. 2019. Learning manifold patch-based representations of man-made shapes. *arXiv preprint arXiv:1906.12337*.
- Teh, C.-H.; and Chin, R. T. 1989. On the detection of dominant points on digital curves. 11(8): 859–872.
- Uy, M. A.; Chang, Y.-Y.; Sung, M.; Goel, P.; Lambourne, J. G.; Birdal, T.; and Guibas, L. J. 2022. Point2cyl: Reverse engineering 3d objects from point clouds to extrusion cylinders. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11850–11860.
- Vahdat, A.; Kreis, K.; and Kautz, J. 2021. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34: 11287–11302.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Vaswani, A. 2017. Attention is All You Need. *arXiv preprint arXiv:1706.03762*.
- Wang, H.; Zhao, M.; Wang, Y.; Quan, W.; and Yan, D.-M. 2024. VQ-CAD: Computer-Aided Design model generation with vector quantized diffusion. *Computer Aided Geometric Design*, 111: 102327.
- Wang, K.; Zheng, J.; and Zhou, Z. 2022. Neural face identification in a 2d wireframe projection of a manifold object. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1622–1631.
- Wang, X.; Xu, Y.; Xu, K.; Tagliasacchi, A.; Zhou, B.; Mahdavi-Amiri, A.; and Zhang, H. 2020. Pie-net: Parametric inference of point cloud edges. *Advances in neural information processing systems*, 33: 20167–20178.
- Willis, K. D.; Jayaraman, P. K.; Lambourne, J. G.; Chu, H.; and Pu, Y. 2021. Engineering sketch generation for computer-aided design. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2105–2114.
- 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.
- 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.; Lambourne, J.; Jayaraman, P.; Wang, Z.; Willis, K.; and Furukawa, Y. 2024. BrepGen: A B-rep Generative Diffusion Model with Structured Latent Geometry. *ACM Transactions on Graphics (TOG)*, 43(4): 1–14.
- 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*.
- Yu, F.; Chen, Q.; Tanveer, M.; Amiri, A. M.; and Zhang, H. 2023. Dualcsg: Learning dual csg trees for general and compact cad modeling. *arXiv preprint arXiv:2301.11497*.
- Yu, F.; Chen, Z.; Li, M.; Sanghi, A.; Shayani, H.; Mahdavi-Amiri, A.; and Zhang, H. 2022. Capri-net: Learning compact cad shapes with adaptive primitive assembly. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11768–11778.
- Zheng, G.; Zhou, X.; Li, X.; Qi, Z.; Shan, Y.; and Li, X. 2023. Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22490–22499.