

From 2D CAD Drawings to 3D Parametric Models: A Vision-Language Approach

Xilin Wang^{1*}, Jia Zheng^{2*}, Yuanchao Hu², Hao Zhu², Qian Yu^{1†}, Zihan Zhou^{2†}

¹Beihang University

²Manycore Tech Inc.

{wang_xilin, qianyu}@buaa.edu.cn, {jiajia, tiancai, hunyuan, shuer}@qunhemail.com

Abstract

In this paper, we present CAD2PROGRAM, a new method for reconstructing 3D parametric models from 2D CAD drawings. Our proposed method is inspired by recent successes in vision-language models (VLMs), and departs from traditional methods which rely on task-specific data representations and/or algorithms. Specifically, on the input side, we simply treat the 2D CAD drawing as a raster image, regardless of its original format, and encode the image with a standard ViT model. We show that such an encoding scheme achieves competitive performance against existing methods that operate on vector-graphics inputs, while imposing substantially fewer restrictions on the 2D drawings. On the output side, our method auto-regressively predicts a general-purpose language describing 3D parametric models in text form. Compared to other sequence modeling methods for CAD which use domain-specific sequence representations with fixed-size slots, our text-based representation is more flexible, and can be easily extended to arbitrary geometric entities and semantic or functional properties. Experimental results on a large-scale dataset of cabinet models demonstrate the effectiveness of our method.

Webpage —

<https://manycore-research.github.io/CAD2Program>

1 Introduction

In computer-aided design (CAD), 2D engineering drawings have been employed as a standard means for describing product designs. With the wide adoption of 2D CAD software (*e.g.*, Autodesk AutoCAD), a significant amount of existing products are presented in the form of engineering drawings today. To manufacture these products, the corresponding 3D objects have to be reconstructed from the 2D drawings.

As shown in Figure 2(a), an engineering drawing typically includes multiple orthographic views of the object.¹ Each orthographic view is a projection of the object onto the plane

*These authors contributed equally.

†Correspondence authors

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

¹Depending on the complexity of the object, a varying number of views may be provided, including a top view, a front view, a side view, and one or more section views.




Figure 1: Illustration of the geometry and annotation layers of a CAD drawing. See text for details.

that is perpendicular to one of the three principal axes. A line of work related to the automatic reconstruction of 3D solid models from orthographic views has existed since the 1970s (Idesawa 1973). However, in the current design and manufacturing industry, human labor is still extensively used to manually realize these 3D object models, which is a time-consuming, error-prone, and unproductive process. We are yet to see the successful application of automatic techniques in commercial CAD software.

To understand the challenges faced by existing automatic techniques, we must take a close look at (i) the engineering drawings and (ii) the 3D models.

Understanding the 2D drawings. An engineering drawing poses a unique challenge as it is a mixture of two types of representations (layers):

- the *geometry layer*, which is the actual object described by its orthographic projections, and
- the *annotation layer*, which includes dimensioning and functional symbols, such as surface types, manufacturing instructions, *etc.*

Figure 1 illustrates the two layers. Some lines constitute the geometry layer of the drawing, as they correspond to the visible edges of the object in the orthographic projection. The other entities, including both text and graphics, form the an-




Figure 2: Problem statement. Given (a) a 2D CAD drawing of a product (e.g., a cabinet), our goal is to reconstruct (b) the 3D model of a product. (c) In CAD software, the 3D model is conventionally built by assembling pre-defined primitive models, where (d) each primitive model is defined by a computer program describing its model ID and a number of parameters.

notation layer. For example, many of the numbers, together with the associated lines, specify the precise sizes (width, depth, and height) of the object parts. A red circle denotes that the shelf is adjustable, whereas the red triangle indicates a door and its opening direction.

Most existing methods for 3D reconstruction from engineering drawings focus on the *geometry layer only*, i.e., to reconstruct 3D B-rep or CSG models by matching the boundary projections across orthographic views (Sakurai and Gossard 1983; Gu, Tang, and Sun 1986; Lequette 1988; Yan, Philip Chen, and Tang 1994; You and Yang 1996; Shin and Shin 1998; Kuo 1998; Liu et al. 2001; Gong et al. 2006a,b). They all assume that a clean geometry layer with exactly three axis-aligned views (namely, front, top, and side views) is provided, and any “irrelevant” annotation has been removed. In practice, this assumption is problematic for at least two reasons: *First*, separating the annotation from the geometry layer is a non-trivial task. As one can see in Figure 1, some lines correspond to the visible edges of the object and belong to the geometry layer, while others belong to the annotation layer (e.g., dimension sets). Thus, the customary text/graphics layer separation does not work here. *Second*, by omitting the annotations, important information about the 3D object, such as dimensions and functional symbols, is ignored during the reconstruction process.

Understanding the 3D models. Now we turn our attention to the output. In CAD software, B-rep is a universal data format, which represents a 3D object as a volume contained in a set of parametric surfaces together with the relationships between them (i.e., the topology). However, human designers rarely create product designs by directly working with these low-level geometric entities. Instead, they make use of pre-defined primitive models (or “parts”) to (i) speed up the design process, and (ii) reduce design errors and ensure that the product can be properly manufactured.

Consider a specific type of product, namely the *cabinet furniture*, for example. Some common primitives may include “base box”, “door”, “drawer”, “fixed shelf”, “ad-




Figure 3: Illustration of the model-specific parameters of a “base box” primitive. N is the number of vertically divided spaces in the box. [NKA, NKB, ...] are the widths of the divided spaces. DBXX indicates the position of the frame, where DBXX=1 means “no frame”, DBXX=2 means “lower frame”, and DBXX=3 means “upper frame”.

justable shelf”, and so on. Figure 2(c) shows an example of decomposing a simple cabinet into the primitives. As one can see in Figure 2(d), each primitive itself is a parametric model. To use a primitive in the product design, one must specify the following:

- the *model ID*, which is the unique identifier of a primitive in the database,
- the *common parameters*, which indicate the general pose and size of the primitive in the 3D space and are model-agnostic, and
- the *model-specific parameters*, which describe possible variations of a specific primitive (see Figure 3 for a visualization).

Since a parametric model can be specified by a list of parameters, a line of recent studies (Nash et al. 2020; Jones et al. 2020; Wu, Xiao, and Zheng 2021; Willis et al. 2021; Ganin et al. 2021; Seff et al. 2022; Guo et al. 2022; Xu et al. 2022; Yu et al. 2023; Jayaraman et al. 2023; Xu et al. 2023) leverage Transformer-based sequence modeling techniques to auto-regressively predict these parameters for CAD model generation. To this end, they define domain-specific sequence representations by collecting all possible parameters and assigning a fixed slot to each parameter in the sequence. One issue with these representations is that, due to the presence of model-specific parameters, the number of slots needed grows with the set of primitives. As a result, these methods typically deal with a small set of (*e.g.*, ≤ 10) primitives.

Our contributions. In this paper, we address the aforementioned challenges as follows.

First, on the input side, instead of extracting and analyzing the graphic elements in the geometry layer as existing methods do, we advocate a more holistic approach in which we simply treat the entire 2D drawing as a raster image and process it with an image encoder (*e.g.*, a ViT). This way, our method not only makes use of information in both geometry and annotation layers for reconstruction, but also eliminates the requirements imposed by existing methods such as a separate geometry layer being available, or exactly three orthographic views (*i.e.*, front, top, and side) are present.

Second, on the output side, we represent 3D parametric models as scripts of a general-purpose language (*e.g.*, Python) and auto-regressively predict the language in text form. Our design choice offers several advantages: (i) it is interpretable and semantically rich; (ii) it can seamlessly integrate new primitives without affecting the length of existing scripts; and (iii) it allows us to leverage strong built-in coding capability of pre-trained large language models.

Finally, with the above input-output setting, we implement our method, CAD2PROGRAM, by directly fine-tuning an open-source vision-language foundation model. In this paper, we choose InternVL (Chen et al. 2024a,b) as the base model. To validate our design choices, we have collected a dataset consisting of 368K cabinet models with 2D engineering drawings. Comprehensive experiments on the dataset demonstrate the effectiveness of our method.

2 Related Work

2.1 3D Reconstruction from Orthographic Views

The problem of 3D reconstructing from orthographic views is a long-standing problem in CAD (Wang and Grinstein 1993). (Idesawa 1973) and (Markowsky and Wesley 1980; Wesley and Markowsky 1981) are the first to present a principled mathematical framework to this problem, and propose to recover 3D vertices, edges, faces, and solids progressively from the input views. Subsequent studies (Sakurai and Gosard 1983; Gu, Tang, and Sun 1986; Lequette 1988; Yan, Philip Chen, and Tang 1994; You and Yang 1996; Shin and Shin 1998; Kuo 1998; Liu et al. 2001; Gong et al. 2006a,b) largely follow the same scheme and focus on (i) improving the computational efficiency of existing algorithms and (ii)

extending the applicability of the framework to handle more complex objects (*e.g.*, quadric surface solids).

Recently, a couple of learning-based approaches emerge as an alternative to the above scheme (Han et al. 2020; Hu et al. 2023). The most relevant work to ours is PlankAssembly (Hu et al. 2023), which also aims to reconstruct 3D parametric models from 2D orthographic views. It employs an encoder-decoder architecture to learn an implicit mapping of geometric entities between the 3D model and 2D views. The main advantage of PlankAssembly is its robustness against errors and missing components in 2D drawings.

However, PlankAssembly has several limitations. *First*, on the input side, it requires exactly three axis-aligned views in vector-graphics format as input, and encodes each entity (*i.e.*, 2D line) in the geometry layer as a token. *Second*, on the output side, it deals with a single type of primitive, namely cuboid-shaped planks with six parameters.

2.2 Sequence Modeling for CAD

With the success of Transformer-based sequence modeling techniques (Vaswani et al. 2017) in the NLP field, a line of work (Nash et al. 2020; Jones et al. 2020; Wu, Xiao, and Zheng 2021; Willis et al. 2021; Ganin et al. 2021; Seff et al. 2022; Guo et al. 2022; Xu et al. 2022; Yu et al. 2023; Jayaraman et al. 2023; Xu et al. 2023) propose to design domain-specific languages (DSL) for the shapes of interest and train deep networks to generate CAD models auto-regressively. Other work (Guo et al. 2022; Ma et al. 2024; Li et al. 2024b; Khan et al. 2024) reconstructs CAD construction sequences from voxels or point clouds.

As discussed before, while most existing methods learn domain-specific sequence representations, we represent 3D parametric models as text scripts using a general-purpose language. As a result, our method can efficiently deal with a large number of primitives and model-specific parameters.

3 Method

In this paper, we focus on cabinet furniture, a common type of parametric models, as our subject of study. We first introduce our text-based shape program in Section 3.1, then show how to train vision-language models to reconstruct 3D cabinet models from 2D drawings in Section 3.2.

3.1 Cabinet Shape Program

As shown in Figure 2, a 3D cabinet model is built by assembling a list of primitive instances $\mathcal{Z} = \{Z_i\}_{i=1}^N$. Each primitive instance can be represented as a tuple $Z_i = \{M_i, B_i, P_i\}$, where M_i is a unique ID of the primitive model in the database, $B_i = \{p_i, s_i, r_i\}$ defines a 3D box comprising all the common parameters including the center position $p_i \in \mathbb{R}^3$, size $s_i \in \mathbb{R}^3$, and 1-D rotation angle $r_i \in \mathbb{R}$ around the vertical axis², and $P_i = \{\text{key}_j = \text{value}_j\}_{j=1}^K$ is a list of model-specific parameters associated with the primitive. Note that $P_i = \emptyset$ if the primitive M_i has no model-specific parameters.

²We do not use full 3D rotation because designers rarely rotate a model along the x - and y -axis.

Listing 1: Python shape program describing the cabinet in Figure 2. Every two lines correspond to a primitive model in Figure 2(c).

```

bbox_0 = Bbox(507, 185, 805, 1014, 370, 50, 0)
model_0 = <model_57761062>()
bbox_1 = Bbox(25, 185, 390, 50, 370, 780, 0)
model_1 = <model_57758898>()
bbox_2 = Bbox(532, 195, 390, 964, 350, 780, 0)
model_2 = <model_115813862>(N=1, NKA=928, DBXX=1, BT=18)
bbox_3 = Bbox(532, 185, 390, 928, 330, 18, 0)
model_3 = <model_57253481>()
bbox_4 = Bbox(291, 11, 390, 478, 18, 776, 0)
model_4 = <model_82289390>(openDirection=0, uCover=18,
    dCover=18, lCover=18, rCover=18)
bbox_5 = Bbox(773, 11, 390, 478, 18, 776, 0)
model_5 = <model_82289390>(openDirection=1, uCover=18,
    dCover=18, lCover=18, rCover=18)

```

In this paper, we follow the convention of CAD and use a coordinate system in which the bounding box of the subject (*i.e.*, cabinet) is aligned to the main axes, with its bottom face aligned with the z -axis, and its front face aligned with the $-y$ -axis. The origin coincides with one of the corners of the box, such that the subject lies in the first octant.

Text-based shape program. A common practice for applying sequence modeling in the CAD field is to design domain-specific language (DSL) for the shape of interest, such as 2D sketches (Willis et al. 2021; Ganin et al. 2021; Seff et al. 2022) and 3D solid models (Wu, Xiao, and Zheng 2021; Guo et al. 2022; Xu et al. 2022; Jayaraman et al. 2023; Xu et al. 2023). In these methods, all the parameters in the DSL are aggregated to form a fixed-length command template C . Then, the shape is represented as a sequence of commands $[C_1, C_2, C_3, \dots]$. For example, DeepCAD (Wu, Xiao, and Zheng 2021) deals with 3 types of geometric entities (*i.e.*, lines, arcs, and circles) and 1 type of operation (*i.e.*, extrusion) for 3D solid generation, resulting in a command template C with 16 parameters.

While a command template can represent arbitrarily complex models in theory, it becomes problematic when applied to our task: due to the presence of model-specific parameters, the length of the command template C grows with the set of primitives. With hundreds or more primitives, the command sequence representing the 3D model would be prohibitively long, with a high percentage of unused slots in each command.

To this end, we find that these limitations can be addressed by directly representing each primitive in text form. In this paper, we choose Python as a proxy language and show an example of our text-based representation in Listing 1. Note that we include both the key and value of each model-specific parameter in the language. This way, we eliminate the need to construct a common command template.

Besides its ability to handle an arbitrary number of primitives, we note that our Python-based shape program offers several other benefits: *First*, in existing methods, the continuous value of a parameter in the command is quantized into discrete tokens by a domain-specific tokenizer, resulting in




Figure 4: An overview of the CAD2PROGRAM model.

an inherent quantization error. By treating these values as text, the quantization error can be avoided. *Second*, it allows us to leverage LLM’s strong capacity to perform Python programming. *Third*, the fact that our program is a sequence of text tokens opens up many potential applications such as text-based editing and visual question answering.

3.2 The CAD2PROGRAM Model

Given the input-output settings, our task can be regarded as an application of vision-language foundation models (VLMs). These models aim to integrate and enhance LLM’s capabilities in processing both visual and textual data. In the past year, significant progress has been made in developing both proprietary commercial models, such as OpenAI’s GPT-4V (OpenAI 2023) and Google’s Gemini-1.5 (Gemini Team 2024), and open-source models including LLaVA (Liu et al. 2024a, 2023, 2024b; Li et al. 2024a), MiniGPT-4 (Zhu et al. 2024), InternVL (Chen et al. 2024a,b), CogVLM (Wang et al. 2024), and many more.

In this paper, we choose Mini-InternVL-1.5-2B for its open-source availability³ and excellent balance between performance and model size. Notably, Mini-InternVL-1.5 supports dynamic resolution by splitting the high-resolution image into up to 12 tiles, significantly enhancing the performance on OCR-related tasks. Also, it is trained on high-quality bilingual dataset and has demonstrated robust capacities in handling multi-modal perception tasks in both English and Chinese. These characteristics make it well-suited for our task.

Figure 4 provides an overview of our CAD2PROGRAM model. It adopts a ViT-MLP-LLM architecture, using InternViT-300M as the vision encoder and InternLM2-1.8B as the language model, respectively. These two models are aligned by a MLP projector. For our task, we use a simple phrase “Reconstruct cabinet from image:” as the text prompt to the language model.

For the output, we note that the model ID is just a randomly generated number and does not carry any semantic meaning. To help the network better align the input visual tokens with the primitives in the database, we create a special token for the model ID. Specifically, we consider two features of the primitive: (i) the model name, and (ii) a snapshot image of the primitive rendered with its default model-specific parameter values from a fixed viewpoint. We use

³<https://huggingface.co/OpenGVLab/Mini-InternVL-Chat-2B-V1-5>




Figure 5: Dataset statistics. **Left:** the number of cabinets w.r.t. the number of primitives per cabinet. **Right:** the number of primitives w.r.t. the number of model-specific parameters per primitive.

Chinese-CLIP (Yang et al. 2023)⁴ with a ViT-H/14 image encoder (Dosovitskiy et al. 2020) and a RoBERTa-wwm-large text encoder (Cui et al. 2021) to extract image and text embeddings, respectively. Then, the special token is obtained by directly concatenating these two embeddings.

Dataset. To train and test our model, we collect a new dataset by accessing a large repository of 3D customized cabinet models on an online interior design platform. On the platform, professional designers utilize 3D modeling software to create 3D cabinet models and the corresponding engineering drawings for real-world production.

When creating the dataset, we restrict the size of the cabinet to be between 0.1m and 4.5m to avoid undersized or oversized models, and also eliminate models with more than 48 primitives. After filtering, our dataset contains 368K cabinet models and 2D engineering drawings, with 373 unique pre-defined primitives. The number of model-specific parameters per primitive ranges from 0 to 8. The total number of model-specific parameters is 702 – at least an order of magnitude larger than the number seen in any command template used in prior work. Some statistics of the dataset are shown in Figure 5. Finally, the dataset is divided into 364K/2K/2K samples for training/validation/testing.

Implementation details. We use the SWIFT (Zhao et al. 2024) framework to train CAD2PROGRAM via supervised full-parameter fine-tuning. We utilize the AdamW optimizer (Loshchilov and Hutter 2017) and a cosine learning rate schedule with a linear warm-up for 1K steps. The peak learning rate is 10^{-5} . The model is trained for about 14K iterations, which takes about 1 day using 64 NVIDIA RTX 4090 GPU devices. The total batch size is set to 128. The length of the token sequence is restricted to 4096.

4 Experiments

4.1 Evaluation Metrics

To evaluate our method, we consider the accuracy of (i) model retrieval (for model ID), (ii) 3D reconstruction (for common parameters), and (iii) parameter estimation (for model-specific parameters). For *3D reconstruction*, we follow PlankAssembly (Hu et al. 2023) to report the precision, recall, and F1 score. Specifically, we use Hungarian matching to match the 3D bounding box of predicted primitives

⁴<https://huggingface.co/OFA-Sys/chinese-clip-vit-huge-patch14>

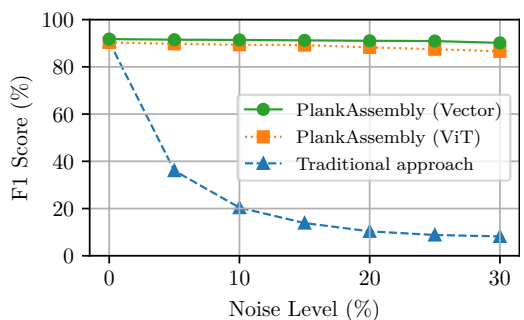


Figure 6: Effect of the image encoder. Results of PlankAssembly (Vector) and the traditional approach are directly taken from (Hu et al. 2023).

with the ground truth. A prediction is regarded as a true positive if its 3D intersection-over-union (IOU) with the ground truth is > 0.5 . For *model retrieval*, we compare each predicted model ID with that of the corresponding ground truth primitive, and report the percentage of matches. For *parameter estimation*, we compute the accuracy on all correctly retrieved models. A successful estimation means that all parameters are correct.

4.2 Experiments on the Input

Effect of the image encoder. We first demonstrate that modern vision models (*i.e.*, ViT) can understand the engineering drawing. For this experiment, we directly modify the PlankAssembly model (Hu et al. 2023), which takes 2D lines from three orthographic views as input, and produces a 3D model consisting of 3D axis-aligned bounding boxes (*i.e.*, planks). To enable PlankAssembly to take image as input, we remove the Transformer encoder (12M) that was used to encode vectorized input, and integrate a TinyViT (Wu et al. 2022)⁵ with comparable model size (21M) to encode the raster image. The decoder then takes the visual tokens generated by TinyViT as input and outputs the 3D model.

To train the modified PlankAssembly model, we use the same code and dataset from its website⁶. For the input raster images, we generate an engineering drawing by arranging the three axis-aligned views on a fixed-size canvas. We also inject noises into the drawings during training and testing, as the original method does.

Figure 6 reports the F1 scores on varying input noise levels. As we can see, the model with the raster image input performs comparably to the model with vectorized input. This suggests that a modern general-purpose vision model is as effective as a domain-specific encoder.

Effect of the annotation layer. One greatest advantage of using an image encoder is that it imposes fewer restrictions on the input as prior methods do, which typically assume that the drawing consists of exactly three orthographic views in vector-graphics format, and that the annotation layer is

⁵https://huggingface.co/timm/tiny_vit_21m_512.dist_in22k_ft_in1k

⁶<https://github.com/manycore-research/PlankAssembly>

Listing 2: Comparison of programs generated by models trained with or without the annotation layer. The corresponding input drawing can be found in Figure 7 (first row, left). Due to the space limit, only the 3D box predictions are shown here. We use underline to indicate numbers that can be directly referred from the dimension sets in the annotation layer. The green and red colors denote the correct and incorrect predictions, respectively. We reorder the predictions by matching them to the ground truth primitives. **Best viewed in color.**

# model trained with annotation layer						# model trained without annotation layer					
2	bbox_10	=	Bbox(1637, 280, 1040, <u>50</u> , 560, <u>2080</u> , 0)	2	bbox_0	=	Bbox(1665, 280, 1040, <u>50</u> , 560, <u>2080</u> , 0)				
3	bbox_9	=	Bbox(1637, 280, 2383, <u>50</u> , 560, <u>606</u> , 0)	3	bbox_1	=	Bbox(1665, 280, 2385, <u>50</u> , 560, <u>610</u> , 0)				
4	bbox_0	=	Bbox(806, 290, 1040, <u>1612</u> , <u>540</u> , <u>2080</u> , 0)	4	bbox_3	=	Bbox(795, 290, 1040, <u>1590</u> , <u>540</u> , <u>2080</u> , 0)				
5	bbox_3	=	Bbox(407.5, 280, 1079, <u>779</u> , 520, 18, 0)	5	bbox_4	=	Bbox(402, 280, 1079, <u>768</u> , 520, 18, 0)				
6	bbox_4	=	Bbox(407.5, 280, 978.5, <u>779</u> , 40, 83, 0)	6	bbox_5	=	Bbox(402, 280, 978.5, <u>768</u> , 40, 83, 0)				
7	bbox_6	=	Bbox(407.5, 280, 1740, <u>779</u> , 520, 18, 0)	7	bbox_6	=	Bbox(402, 280, 1740, <u>768</u> , 520, 18, 0)				
8	bbox_5	=	Bbox(407.5, 280, 1409, <u>779</u> , 520, 18, 0)	8	bbox_7	=	Bbox(402, 280, 1409, <u>768</u> , 520, 18, 0)				
9	bbox_1	=	Bbox(1204.5, 280, 1575, <u>779</u> , 520, 18, 0)	9	bbox_8	=	Bbox(1188, 280, 1585, <u>768</u> , 520, 18, 0)				
10	bbox_2	=	Bbox(1204.5, 280, 1474.5, <u>779</u> , 40, 83, 0)	10	bbox_9	=	Bbox(1188, 280, 1484.5, <u>768</u> , 40, 83, 0)				
11	bbox_7	=	Bbox(806, 290, 2383, <u>1612</u> , <u>540</u> , <u>606</u> , 0)	11	bbox_10	=	Bbox(795, 290, 2385, <u>1590</u> , <u>540</u> , <u>610</u> , 0)				
12	bbox_8	=	Bbox(831, 280, 2746, <u>1662</u> , 560, <u>120</u> , 0)	12	bbox_2	=	Bbox(820, 280, 2750, <u>1640</u> , 560, <u>120</u> , 0)				

input layers		retrieval	reconstruction		param.
geometry	annotation	acc.	prec.	rec.	F1 acc.
✓		85.42	63.00	62.48	62.65 81.94
✓	✓	93.80	83.10	82.56	82.76 97.21

Table 1: Effect of the annotation layer.

removed. In this experiment, we investigate the impact of the annotation layer on the reconstruction task.

For this experiment, we use our new dataset because the PlankAssembly dataset does not provide any annotation. We compare the performance of CAD2PROGRAM model trained with two different inputs: (i) geometry layer only, and (ii) both geometry and annotation layers.

As shown in Table 1, the model taking the original drawing (with both layers) as input achieves significantly higher accuracies across all metrics. Notably, this is in contrast to traditional practice which treats the annotation layer as nuisances that could hurt the performance of a reconstruction method. To gain insight into the performance gain, we show example programs generated by these two models in Listing 2. As one can see, the dimension sets in the annotation provide direct reference w.r.t. the size of the primitives, which may be exploited by a vision-language model. Without the annotation, the model must infer the sizes based on the geometric entities, resulting in less accurate predictions.

4.3 Experiments on the Output

Next, we study the performance of our model with different output representations. To verify the effectiveness of the proposed text-based shape program, we compare it with a domain-specific sequence representation. As discussed in Section 3.1, a method using domain-specific sequence representation with a command template lacks the flexibility to predict model-specific parameters. Thus, for this experiment, we compare CAD2PROGRAM to two variants which predict the model ID and common parameters only:

- The *first variant* outputs domain-specific sequence rep-

resentations. To this end, we create a command template that includes the model ID and all common parameters. For common parameters, we quantize each position and size parameter into 1500 bins with a resolution of 3mm, and the rotation angle into 4 bins. For model ID, we use the same special token as described in Section 3.2. This variant has a standard encoder-decoder architecture, including InternViT-300M as the encoder and a 6-layer Transformer decoder (similar to the decoder used in PlankAssembly (Hu et al. 2023)).

- The *second variant* has the same architecture as CAD2PROGRAM, but is trained to predict a Python-based shape program describing the model ID and common parameters only.

The results are shown in Table 2. We observe that all three models show comparable performance on both model retrieval and 3D reconstruction. This is significant because it suggests that using free-form text as output is as effective as employing a domain-specific sequence representation as in prior work (e.g., PlankAssembly). Moreover, our CAD2PROGRAM model with text-based shape program offers greater flexibility in designing the output sequence and performs well across all metrics, i.e., 3D reconstruction, model retrieval, and model-specific parameter estimation.

Figure 7 visualizes some 3D parametric reconstruction results of CAD2PROGRAM. As one can see, our method successfully reconstructs a wide variety of cabinet models. Note that the input 2D drawings here would cause major problems for previous methods (e.g., PlankAssembly) for several reasons. *First*, they contain a varying number of orthographic views. For example, in Figure 7 (first row, left), the drawing does not show the side view (instead, a section view is provided to better reveal the inside design), whereas in Figure 7 (third row, right), five views are shown. *Second*, the different views in a drawing may not be aligned to each other. *Third*, both the geometry layer and annotation layer are present.

In the last row of Figure 7, we illustrate some common artifacts in the models reconstructed by our method. On the left, we show a case where our method produces a model consisting of two parts, with a small gap in between (high-

format	model ID	output		retrieval acc.	reconstruction			param. acc.
		common param.	model-specific param.		prec.	rec.	F1	
command	✓	✓		93.36	82.46	81.86	82.07	n/a
text	✓	✓		93.84	82.94	82.51	82.65	n/a
text	✓	✓	✓	93.80	83.10	82.56	82.76	97.21

Table 2: Experimental results on output representation.

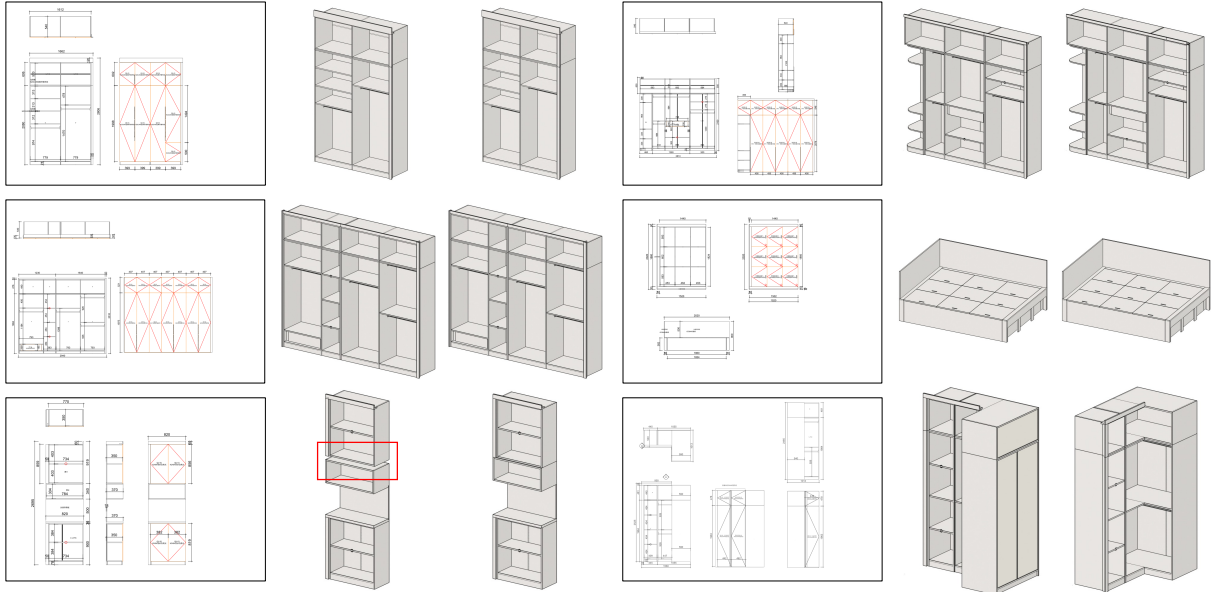


Figure 7: Qualitative results. For each case, we show **from left to right** the input drawing, the model reconstructed by CAD2PROGRAM, and the ground truth. To reveal the inside structure, we manually delete the doors in the visualization.

proxy language	retrieval		reconstruction		param. acc.
	acc.	prec.	rec.	F1	
YAML	93.12	83.18	83.07	83.05	97.09
Python	93.80	83.10	82.56	82.76	97.21

Table 3: Ablation on proxy language.

lighted by the red box). This suggests that the predicted 3D positions (*i.e.*, p_i) may not be accurate enough. Note that, unlike the sizes (*i.e.*, s_i) which can often be directly referred from the annotations, 3D positions need to be inferred, sometimes through arithmetic calculations. This remains a challenging task for LLM. On the right, due to ambiguity in the coordinate system used, the L-shape cabinet is reconstructed with a different pose w.r.t. the ground truth.

Ablation on the proxy language. To examine the impact of our choice of proxy language, we have conducted an experiment in which we replace Python with YAML, a popular data serialization language. Compared to Python, the YAML script is generally longer. Specifically, the average script lengths for 3D models in our dataset are 697 and 1208 tokens for Python and YAML, respectively.

As shown in Table 3, the models trained with both languages yield very similar results, indicating that the choice of proxy language does not significantly affect the model performance for our task.

5 Conclusion

In this paper, we present a vision-language approach to 3D parametric model reconstruction from 2D CAD drawings. Our key observations are: (i) a modern vision encoder (*e.g.*, ViT) is effective in understanding CAD drawings, and (ii) using text-based sequence representation provides greater flexibility in incorporating a large number of primitives with arbitrary geometric, semantic, and functional properties.

The studies presented in this paper are limited to cabinet furniture only. However, we point out that our method is general and can be applied to other types of CAD models – if a large-scale dataset for the domain of interest is available. Meanwhile, the use of vision-language foundation models opens up many directions for future research. A direction of particular interest is visual question answering in the context of 2D CAD drawings, a critical capacity for developing autonomous agents for design and manufacturing.

Acknowledgements

This work was done when Xilin Wang was an intern at Manycore Tech Inc. This work was supported by the National Science and Technology Major Project (No. 2022ZD0117800), Young Elite Scientists Sponsorship Program by CAST, and the Fundamental Research Funds for the Central Universities.

References

- Chen, Z.; Wang, W.; Tian, H.; Ye, S.; Gao, Z.; Cui, E.; Tong, W.; Hu, K.; Luo, J.; Ma, Z.; Ma, J.; Wang, J.; Dong, X.; Yan, H.; Guo, H.; He, C.; Shi, B.; Jin, Z.; Xu, C.; Wang, B.; Wei, X.; Li, W.; Zhang, W.; Zhang, B.; Cai, P.; Wen, L.; Yan, X.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2024a. How Far Are We to GPT-4V? Closing the Gap to Commercial Multimodal Models with Open-Source Suites. arXiv:2404.16821.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; Li, B.; Luo, P.; Lu, T.; Qiao, Y.; and Dai, J. 2024b. InternVL: Scaling up Vision Foundation Models and Aligning for Generic Visual-Linguistic Tasks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 24185–24198.
- Cui, Y.; Che, W.; Liu, T.; Qin, B.; and Yang, Z. 2021. Pre-Training With Whole Word Masking for Chinese BERT. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 29: 3504–3514.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Hounsby, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *Int. Conf. Learn. Represent.*
- Ganin, Y.; Bartunov, S.; Li, Y.; Keller, E.; and Saliceti, S. 2021. Computer-Aided Design as Language. In *Adv. Neural Inform. Process. Syst.*, 5885–5897.
- Gemini Team. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv:2403.05530.
- Gong, J.; Zhang, G.; Zhang, H.; and Sun, J. 2006a. Reconstruction of 3D curvilinear wire-frame from three orthographic views. *Comput. Graph.*, 30(2): 213–224.
- Gong, J.; Zhang, H.; Zhang, G.; and Sun, J. 2006b. Solid reconstruction using recognition of quadric surfaces from orthographic views. *Comput. Aided Des.*, 38(8): 821–835.
- Gu, K.; Tang, Z.; and Sun, J. 1986. Reconstruction of 3D Objects from Orthographic Projections. *Comput. Graph. Forum*, 5(4): 317–323.
- Guo, H.; Liu, S.; Pan, H.; Liu, Y.; Tong, X.; and Guo, B. 2022. ComplexGen: CAD reconstruction by B-rep chain complex generation. *ACM Trans. Graph.*, 41(4): 129:1–129:18.
- Han, W.; Xiang, S.; Liu, C.; Wang, R.; and Feng, C. 2020. SPARE3D: A Dataset for SPATial REasoning on Three-View Line Drawings. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 14678–14687.
- Hu, W.; Zheng, J.; Zhang, Z.; Yuan, X.; Yin, J.; and Zhou, Z. 2023. PlankAssembly: Robust 3D Reconstruction from Three Orthographic Views with Learnt Shape Programs. In *IEEE Int. Conf. Comput. Vis.*, 18495–18505.
- Idesawa, M. 1973. A System to Generate a Solid Figure from Three View. *Bulletin of the JSME*, 16(92): 216–225.
- Jayaraman, P. K.; Lambourne, J. G.; Desai, N.; Willis, K. D. D.; Sanghi, A.; and Morris, N. J. W. 2023. SolidGen: An Autoregressive Model for Direct B-rep Synthesis. *Trans. Mach. Learn. Res.*
- Jones, R. K.; Barton, T.; Xu, X.; Wang, K.; Jiang, E.; Guerrero, P.; Mitra, N. J.; and Ritchie, D. 2020. ShapeAssembly: learning to generate programs for 3D shape structure synthesis. *ACM Trans. Graph.*, 39(6): 234:1–234:20.
- Khan, M. S.; Dupont, E.; Ali, S. A.; Cherenkova, K.; Kacem, A.; and Aouada, D. 2024. CAD-SIGNet: CAD Language Inference from Point Clouds using Layer-wise Sketch Instance Guided Attention. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4713–4722.
- Kuo, M.-H. 1998. Reconstruction of quadric surface solids from three-view engineering drawings. *Comput. Aided Des.*, 30(7): 517–527.
- Lequette, R. 1988. Automatic construction of curvilinear solids from wireframe views. *Comput. Aided Des.*, 20(4): 171–180.
- Li, B.; Zhang, K.; Zhang, H.; Guo, D.; Zhang, R.; Li, F.; Zhang, Y.; Liu, Z.; and Li, C. 2024a. LLaVA-NeXT: Stronger LLMs Supercharge Multimodal Capabilities in the Wild. <https://llava-vl.github.io/blog/2024-05-10-llava-next-stronger-llms/>. Accessed: 2024-12-16.
- Li, P.; Guo, J.; Li, H.; Benes, B.; and Yan, D.-M. 2024b. SfmCAD: Unsupervised CAD Reconstruction by Learning Sketch-based Feature Modeling Operations. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4671–4680.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved Baselines with Visual Instruction Tuning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 26296–26306.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024b. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge. <https://llava-vl.github.io/blog/2024-01-30-llava-next/>. Accessed: 2024-12-16.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual Instruction Tuning. In *Adv. Neural Inform. Process. Syst.*, 34892–34916.
- Liu, S.; Hu, S.; Chen, Y.; and Sun, J. 2001. Reconstruction of curved solids from engineering drawings. *Comput. Aided Des.*, 33(14): 1059–1072.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled Weight Decay Regularization. In *Int. Conf. Learn. Represent.*
- Ma, W.; Chen, S.; Lou, Y.; Li, X.; and Zhou, X. 2024. Draw Step by Step: Reconstructing CAD Construction Sequences from Point Clouds via Multimodal Diffusion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 27154–27163.
- Markowsky, G.; and Wesley, M. A. 1980. Fleshing Out Wire Frames. *IBM J. Res. Dev.*, 24(5): 582–597.

- Nash, C.; Ganin, Y.; Eslami, S. M. A.; and Battaglia, P. W. 2020. PolyGen: An Autoregressive Generative Model of 3D Meshes. In *Int. Conf. Mach. Learn.*, 7220–7229.
- OpenAI. 2023. GPT-4V(ision) System Card. https://cdn.openai.com/papers/GPTV_System_Card.pdf. Accessed: 2024-12-16.
- Sakurai, H.; and Gossard, D. C. 1983. Solid Model Input through Orthographic Views. In *ACM SIGGRAPH*, 243–252.
- Seff, A.; Zhou, W.; Richardson, N.; and Adams, R. P. 2022. Vitruvion: A Generative Model of Parametric CAD Sketches. In *Int. Conf. Learn. Represent.*
- Shin, B.; and Shin, Y. 1998. Fast 3D Solid Model Reconstruction from Orthographic Views. *Comput. Aided Des.*, 30(1): 63–76.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is All you Need. In *Adv. Neural Inform. Process. Syst.*, 5998–6008.
- Wang, W.; and Grinstein, G. G. 1993. A Survey of 3D Solid Reconstruction from 2D Projection Line Drawings. *Comput. Graph. Forum*, 12(2): 137–158.
- Wang, W.; Lv, Q.; Yu, W.; Hong, W.; Qi, J.; Wang, Y.; Ji, J.; Yang, Z.; Zhao, L.; Song, X.; Xu, J.; Xu, B.; Li, J.; Dong, Y.; Ding, M.; and Tang, J. 2024. CogVLM: Visual Expert for Pretrained Language Models. arXiv:2311.03079.
- Wesley, M. A.; and Markowsky, G. 1981. Fleshing Out Projections. *IBM J. Res. Dev.*, 25(6): 934–953.
- Willis, K. D. D.; Jayaraman, P. K.; Lambourne, J. G.; Chu, H.; and Pu, Y. 2021. Engineering Sketch Generation for Computer-Aided Design. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2105–2114.
- Wu, K.; Zhang, J.; Peng, H.; Liu, M.; Xiao, B.; Fu, J.; and Yuan, L. 2022. TinyViT: Fast Pretraining Distillation for Small Vision Transformers. In *Eur. Conf. Comput. Vis.*, 68–85.
- Wu, R.; Xiao, C.; and Zheng, C. 2021. DeepCAD: A Deep Generative Network for Computer-Aided Design Models. *IEEE Int. Conf. Comput. Vis.*, 6772–6782.
- Xu, X.; Jayaraman, P. K.; Lambourne, J. G.; Willis, K. D. D.; and Furukawa, Y. 2023. Hierarchical Neural Coding for Controllable CAD Model Generation. In *Int. Conf. Mach. Learn.*, 38443–38461.
- Xu, X.; Willis, K. D. D.; Lambourne, J. G.; Cheng, C.; Jayaraman, P. K.; and Furukawa, Y. 2022. SkexGen: Autoregressive Generation of CAD Construction Sequences with Disentangled Codebooks. In *Int. Conf. Mach. Learn.*, 24698–24724.
- Yan, Q.; Philip Chen, C. L.; and Tang, Z. 1994. Efficient algorithm for the reconstruction of 3D objects from orthographic projections. *Comput. Aided Des.*, 26(9): 699–717.
- Yang, A.; Pan, J.; Lin, J.; Men, R.; Zhang, Y.; Zhou, J.; and Zhou, C. 2023. Chinese CLIP: Contrastive Vision-Language Pretraining in Chinese. arXiv:2211.01335.
- You, C.; and Yang, S. 1996. Reconstruction of curvilinear manifold objects from orthographic views. *Comput. Graph.*, 20(2): 275–293.
- Yu, F.; Chen, Q.; Tanveer, M.; Mahdavi-Amiri, A.; and Zhang, H. 2023. D²CSG: Unsupervised Learning of Compact CSG Trees with Dual Complements and Dropouts. In *Adv. Neural Inform. Process. Syst.*, 22807–22819.
- Zhao, Y.; Huang, J.; Hu, J.; Zhang, D.; Jiang, Z.; Wu, Z.; Ai, B.; Wang, A.; Zhou, W.; and Chen, Y. 2024. SWIFT: A Scalable lightWeight Infrastructure for Fine-Tuning. arXiv:2408.05517.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2024. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. In *Int. Conf. Learn. Represent.*