

TransCAD: A Hierarchical Transformer for CAD Sequence Inference from Point Clouds

Elona Dupont¹, Kseniya Cherenkova^{1,2}, Dimitrios Mallis¹, Gleb Gusev², Anis Kacem¹, and Djamila Aouada¹

¹ SnT, University of Luxembourg, Luxembourg

² Artec 3D, Luxembourg

Abstract. 3D reverse engineering, in which a CAD model is inferred given a 3D scan of a physical object, is a research direction that offers many promising practical applications. This paper proposes *TransCAD*, an end-to-end transformer-based architecture that predicts the CAD sequence from a point cloud. *TransCAD* leverages the structure of CAD sequences by using a hierarchical learning strategy. A loop refiner is also introduced to regress sketch primitive parameters. Rigorous experimentation on the DeepCAD [43] and Fusion360 [41] datasets show that *TransCAD* achieves state-of-the-art results. The result analysis is supported with a proposed metric for CAD sequence, the *mean Average Precision of CAD Sequence*, that addresses the limitations of existing metrics.

Keywords: Computer-Aided Design · CAD Reverse Engineering · Feature-based modelling

1 Introduction

Practically every object encountered in daily life originates from a Computer-Aided Design (CAD), highlighting the fundamental role of CAD in industrial manufacturing processes. Currently, the dominant paradigm for CAD design is *feature-based modelling* [48]. It allows the creation and manipulation of 3D models through a series of *features*, individual elements or operations (holes, slots, fillets, *etc.*), that modify the geometry of a CAD model. The process is typically initiated with the design of *planar sketches*, *i.e.* collection of *loops* composed of 2D curves, followed by a CAD operation (extrusion, revolution, *etc.*) that expands sketches into a 3D solid model. The final model is represented by the sequence of these CAD sketches and operations. *Feature-based* modelling has been widely adopted, as it enables intuitive design alterations and seamless CAD software integration, making it essential for an iterative development of complex designs. The recent availability of large CAD model datasets, such as ABC [18] and Fusion360 [41], has sparked significant interest in developing learning-based approaches for feature-based modelling. Recent efforts have been focused on deep generative modelling [33, 43, 45, 47], where large transformer-based networks are trained to create new CAD models or to automatically complete partial designs

via autoregressive inference. While this research direction offers a lot of potential practical applications for CAD software integration, far less attention has been put to *reverse engineering*. Feature-based reverse engineering emerges as a real-world application addressing the need to automatically replicate physical objects as CAD models. Recovery of CAD design is facilitated by the acquisition of a point cloud or triangular mesh of a physical object scanned using commercial 3D sensors.

Some existing reverse engineering approaches investigate the recovery of alternative CAD model representations like Constructive Solid Geometry [15] (CSG) or Boundary-Representation (B-Rep) [8, 14, 19]. Other methods tackle feature-based reverse engineering and predict implicit representations of sketches and CAD operations from point clouds [23, 36]. Nevertheless, such approaches do not allow for seamless integration into CAD software and often require post-processing (*e.g.* parametric curve fitting). To address these limitations, models capable of learning explicit CAD sequence of parametric sketches and operations from point clouds are needed. This can be enabled within a generative learning framework as in [43]. In that work, an auto-encoder reconstructing CAD parametric sequences is proposed and the latent representation is used for generating novel CAD sequences. An extension for reverse engineering was proposed by replacing the CAD sequence encoder with a point cloud encoder trained to map point clouds to the latent representations. The main limitation of the above is the predefined latent space that cannot adapt to the variations present in real-world point clouds. This disconnection can cause the model to generalize poorly to unseen inputs, especially those with noise or irregularities that are present in 3D scans and that are not well-represented in the training data.

To that end, we propose *TransCAD*, a novel end-to-end trainable and single-stage hierarchical network for feature-based explicit CAD sequence reverse engineering from point clouds. Our network is hierarchical in the sense that it employs a two-tiered decoding process. Initially, a primary CAD sequence embedding is decoded, encapsulating high-level features of the design, that are then processed by secondary decoders, one dedicated to loop parameters and another to CAD operations. Each decoder specializes in a certain input, enabling a nuanced and precise recovery of CAD parameters. The decomposition of learned representations matches the decomposition inherent in the actual feature-based design process of conceptualizing a 3D model through distinct loop and operation steps. Moreover, *TransCAD* does not predict sketch primitive types explicitly as in [43]; instead, we employ a unified primitive representation where types are determined solely by coordinates. Our formulation narrows the learning space by eliminating syntactically incorrect predictions and facilitates a seamless transition between different primitive types. Additionally, it allows for a cascaded parameter refining that further enhances model performance.

Another focus of this work is the evaluation of parametric CAD sequence. We identify several limitations of the existing evaluation framework used by [27, 43] and suggest a suitable metric for CAD sequence similarity based on mean average precision, computed in the unquantized parametric space.

Contributions: In summary our contributions are the following:

1. We propose *TransCAD*, a novel hierarchical architecture for feature-based reverse engineering. Our model is single-stage and end-to-end trainable. *TransCAD* allows for a compact CAD sequence representation that does not include categorical types and enables cascaded coordinate refinement.
2. We identify several limitations of the existing evaluation framework for feature-based reverse engineering and propose a new evaluation metric framework to ensure fair comparison among diverse network architectures.
3. Our model surpasses the performance of recent generative-based approaches while also bridging the gap to real-world applications by exhibiting robustness to perturbed point clouds.

Paper Organization: The rest of the paper is organized as follows. Section 2 reviews related works. Section 3 formulates the problem of feature-based CAD reverse-engineering. The proposed *TransCAD* is described in Section 4. Discussion on the current evaluation framework and suggested extension is introduced in Section 5. An experimental validation of the proposed network is provided in Section 6. Finally, conclusions are given in Section 7.

2 Related Works

Generative Models for CAD: The advent of large-scale 3D shape datasets [4, 18, 24, 41], combined with the significant progress for generative models in vision [5, 10, 12, 16], has sparked interest in the generation of 3D shapes. Existing methods have been proposed for various 3D representations, including point clouds [1, 2, 49], 3D meshes [29, 38], voxel grids [21, 42], and signed distance functions [6, 44]. This work focuses on CAD model generation, which compared to the above is parametric and directly editable in CAD software. A line of work explores the generation of the Boundary-Representation (B-Rep), a collection of parametric surfaces connected via a structured topological graph. SolidGen [13] considers B-Rep synthesis based on transformers and two-level pointer networks. BrepGen [46] represents a B-Rep via a fixed tree of latent geometry representations that can be generated by a diffusion model. Feature-based CAD generation has also been recently explored. Most relevant to our work is DeepCAD [43], a non-autoregressive generative model capable of synthesizing novel CAD sequences based on a transformer auto-encoder architecture. In [45, 47] the authors also follow autoregressive strategies. HNC [45] uses a hierarchical model based on high-level concepts and a code tree for CAD model generation and auto-completion. Similarly, in SkexGen [47] a transformer architecture is used to generate CAD models in the sketch-extrude format by encoding the topology and geometry using different codebooks. All the aforementioned works are oriented around 3D shape generation and either do not address the reverse-engineering task or address it via adaptation of generative modelling leading to suboptimal performance.

CAD Reverse Engineering: Reverse engineering is a well-studied problem with a substantial research effort directed towards predicting geometric features of CAD models, by analyzing the corresponding point clouds. Parametric fitting techniques infer the parameterization of edges [7, 25, 34, 39, 51, 52] and surfaces [11, 22]. Various attributes of the B-Rep and CAD operations are recovered from 3D scans in [28]. CADOps-Net [8] recovers 2D sketches from faces segmented into their CAD operation steps. Reasoning about a CAD model via properties discovered by parametric fitting offers insights solely into its end-state, without considering the sequential CAD design process intrinsic to feature-based modeling.

A step closer to CAD reverse engineering, a line of work explores the reconstruction of a point cloud into Constructive Solid Geometry (CSG) [9, 15, 50], a modelling technique that uses boolean operations to combine primitives into 3D models. Point2Cyl [36], on the other hand, predicts extrusion cylinders given a point cloud, but requires user input to combine cylinders. SECAD-Net [23] and ExtrudeNet [33] use a self-supervised learning strategy to recover CAD sequences in the form of implicit representations given voxels and point clouds, respectively. In contrast to feature-based modelling, 3D representations produced by these methods (CSG, extrusion cylinders, *etc.*) have limited compatibility with modern CAD software workflows. The authors in [20] learn sketch-extrude sequences conditioned on a voxel input, however, the model relies on strong data priors and is limited to predefined extrusion combinations.

Closer to our work is DeepCAD [43] and subsequent MultiCAD [27]. Even though DeepCAD [43] proposes a non-autoregressive generative framework for feature-based CAD, authors explore further conditioning on input point clouds. Taking a similar direction, MultiCAD [27] proposes a two-stage multimodal contrastive learning strategy of both point clouds and CAD sequences. The two aforementioned methods opt for separate stage learning for point clouds and CAD sequences. Concurrent to our work, the autoregressive strategy in [17] and the multimodal diffusion based approach in [26] attempt to solve the point cloud to CAD sequence problem. To our knowledge *TransCAD* is the first non-autoregressive single-stage architecture for feature-based reverse engineering.

3 Problem Statement

A CAD model $\mathbf{C} \in \mathcal{C}$ is constructed in a sequence of construction steps. Each step can be seen as a 2D parametric *sketch* $\mathbf{s} \in \mathcal{S}$ (*e.g.* set of lines, arcs, *etc.*) followed by a *CAD operation* $\mathbf{o} \in \mathcal{O}$ (*e.g.* extrusion, revolution, *etc.*) [43, 47]. Here, \mathcal{C} is the set of all possible CAD models, \mathcal{S} and \mathcal{O} represent the sets of possible CAD sketches and operations, respectively. CAD models constructed exclusively from the *extrusion* operation type are considered in this work. Extrusion $\mathbf{e} \in \mathcal{E}$, where \mathcal{E} denotes the set of possible extrusions, is the most common operation and enables the description of a wide range of CAD models [43, 45, 47]. *TransCAD* aims at learning how to predict the sequence of CAD construction steps from an input point cloud. Formally, given a point cloud $\mathbf{P} = [\mathbf{p}_1, \dots, \mathbf{p}_n] \in \mathbb{R}^{n \times 3}$, where

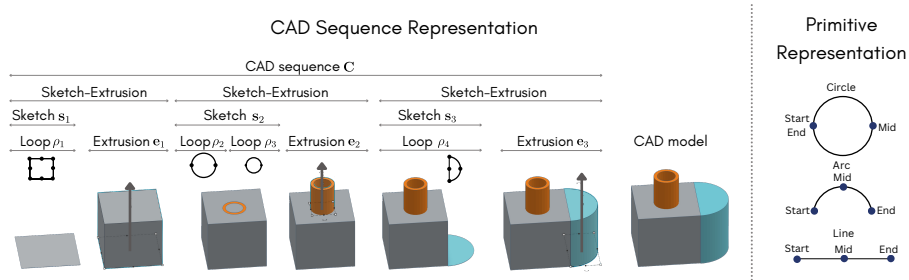


Fig. 1: The sequential process of CAD modeling. A CAD sequence \mathbf{C} can be decomposed into a hierarchical structure. The highest conceptual level is a sequence of sketch \mathbf{s} and extrusion \mathbf{e} . A sketch can be made of one or more loops ρ . Each loop can be further decomposed into loop primitives, circle, arc and line. Each loop primitive can be described by a fixed number of parameters as shown on the right panel.

$\mathbf{p}_i = [x_i, y_i, z_i]$ denotes the 3D coordinates of the point i and n the number of points, the objective of *TransCAD* is to learn the mapping

$$\begin{aligned} \Phi &: \mathbb{R}^{n \times 3} \rightarrow \mathcal{C}, \\ \Phi(\mathbf{P}) &= \{\mathbf{s}_l, \mathbf{e}_l\}_{l=1}^L, \end{aligned} \tag{1}$$

where L denotes the length of the CAD sequence. In what follows, the proposed formulations of sketches and extrusions are described.

CAD Sketch and Extrusion Formulation: The proposed formulations for sketch and extrusion steps are inspired by [43, 47]. A sketch \mathbf{s} is composed of one or more *loops* (see left panel of Fig. 1). Each loop $\{\rho_j\}_{j=1}^{L_\rho}$, where L_ρ , denoting the number of loops, consists of one primitive (*i.e.* circle) or a combination of primitives (*i.e.* lines and arcs). In contrast to [43] which specifies the type of primitives in their representation, the proposed primitive representation is type-agnostic (see right panel of Fig. 1). In particular, each primitive δ is represented by three 2D coordinates of start, mid, and end points, *i.e.* $\delta = [(x_{start}, y_{start}), (x_{mid}, y_{mid}), (x_{end}, y_{end})] \in \mathbb{R}^6$. This representation has the advantage that the type of primitive can be deduced from the configurations of the points, hence reducing the search space and facilitating the transition between different types during training. In practice, the mid point of a line is replaced by a dummy value.

As in [43], we ensure that the loops are always closed by using the end point of a primitive as the start point of the next one. Further, a similar to [43] quantization is considered to reduce the parameters search space. As a result, a loop of n_p primitives $\rho_j \in \mathbb{R}^{6 \times n_p}$ is considered as a quantized representation $\rho_j^* \in \llbracket 0..d_q \rrbracket^{6 \times n_p}$, where d_q denotes the quantization interval. As for extrusion, similarly to [43], a quantized representation $\mathbf{e}_j^* \in \llbracket 0..d_q \rrbracket^{11}$ is considered to represent the sketch plane/scale and extrusion type/distances. Note that $\mathbf{e}^* \in \llbracket 0..d_q \rrbracket^{11 \times L_e}$ and $\rho^* \in \llbracket 0..d_q \rrbracket^{6 \times n_p \times L_\rho}$ will be used in the following to




Fig. 2: *TransCAD* model architecture. *TransCAD* is a hierarchical network composed of the following components: a point cloud encoder, a loop-extrusion decoder that predicts a high-level sequence which is then decoded by a loop decoder and an extrusion decoder. The predicted quantized loop parameters are then corrected by a loop refiner.

denote sequences of quantized extrusions and loops, respectively. Here, L_{ρ} and L_e denote the length of loop and extrusion sequences, respectively.

4 Hierarchical CAD Sequence Learning from Point Clouds

TransCAD non-autoregressively learns to predict a CAD sequence from an input point cloud in the format described in Section 3. First, the point cloud is encoded into point features using a standard point cloud encoder. In order to facilitate the learning of CAD sequences, a hierarchical CAD sequence decoding is proposed. In particular, a high-level sequence of embedding corresponding to the loop and extrusion steps is learned. Those embedding are then fed to either a loop or extrusion decoder based on predicted type to learn the loop and extrusion parameters. Finally, the predicted loop parameters are further refined using the actual unquantized loop parameters (ground truth). The overall model architecture is depicted in Fig. 2 and the different components are described below.

4.1 Point Cloud Encoder

The point cloud encoder Φ_p consists of 4 layers of PointNet++ [32] and operates on an input point cloud \mathbf{P} . It outputs per-point features encoding local neighborhood information, $\mathbf{F}_p = [\mathbf{f}_p^1 \dots \mathbf{f}_p^n] \in \mathbb{R}^{n \times d_p}$, that encode local neighborhood information, d_p denotes the dimension of the features. Note that point normals of \mathbf{P} are estimated using [35] and are provided as input to Φ_p along with its 3D coordinates.

4.2 Loop-Extrusion Decoder

The main objective of the loop-extrusion decoder $\Phi_{\rho,e}$ is to learn a high-level sequence of embedding $\mathbf{F}_{\rho,e} = [\mathbf{f}_{\rho,e}^1, \dots, \mathbf{f}_{\rho,e}^{L_{\rho,e}}] \in \mathbb{R}^{d_z \times L_{\rho,e}}$ corresponding to loops and extrusions from the point cloud features \mathbf{F}_p . Here, d_z and $L_{\rho,e}$ denote the embedding dimension and the length of the sequence, respectively.

The decoder $\Phi_{\rho,e}$ is composed of multi-head transformer-based blocks [37]. In the first block, learned constant embedding $\mathbf{F}_c \in \mathbb{R}^{d_z \times L_{\rho,e}}$ undergo a self-attention operation [37] and the resulting representation cross-attends to the point cloud features \mathbf{F}_p to produce loop extrusion embedding for the first block $\mathbf{F}_{\rho,e}^1 \in \mathbb{R}^{d_z \times L_{\rho,e}}$ as follows,

$$\mathbf{F}_{\rho,e}^1 = \text{CA}(\text{SA}(\mathbf{F}_c), \mathbf{F}_p), \quad (2)$$

where $\text{SA}(\cdot)$ and $\text{CA}(\cdot, \cdot)$ denote the self and cross attention operators [37], respectively. The same self and cross attention operations are conducted in the subsequent blocks by feeding the output of each block as input to the next one,

$$\mathbf{F}_{\rho,e}^b = \text{CA}(\text{SA}(\mathbf{F}_{b-1}), \mathbf{F}_p), \quad (3)$$

yielding the final sequence of embedding $\mathbf{F}_{\rho,e}$ after the last block. In order to ensure that each element $\mathbf{f}_{\rho,e}^i$ in the sequence embedding $\mathbf{F}_{\rho,e}$ corresponds to the right type (*i.e.* loop, extrusion, or end of sequence), a 3 layer MLP followed by **softmax** that operates on each $\mathbf{f}_{\rho,e}^i$ and predicts its type is introduced. A cross-entropy loss, $\mathcal{L}_{\rho,e}$, is computed between the predicted and ground truth types to supervise the learning of $\mathbf{F}_{\rho,e}$. Note that the loop-extrusion decoder is solely used to obtain a high-level sequence of loop and extrusion embedding. These embedding can then be decoded through either a loop decoder or an extrusion decoder to obtain their parameters. At training time, the ground truth type labels are used to identify which decoder should be used for each embedding, while at inference time the predicted types are used. The identification of loop and extrusion types results into separate loop embedding $\mathbf{F}_\rho \in \mathbb{R}^{d_z \times L_\rho}$ and extrusion embedding $\mathbf{F}_e \in \mathbb{R}^{d_z \times L_e}$ by splitting $\mathbf{F}_{\rho,e}$ according to loop and extrusion types.

4.3 Loop and Extrusion Parametrization

After obtaining the representation and the type of loop and extrusion steps, the parameters of both loops and extrusions are decoded from these representations using separate decoders.

Extrusion Decoder: As mentioned in Section 3, the extrusion sequence is described by a sequence of 11 quantized parameters $\mathbf{e}^* \in \llbracket 0..d_q \rrbracket^{11 \times L_e}$. In order to obtain these parameters from the extrusion sequence embedding \mathbf{F}_e , an extrusion decoder Φ_e consisting of 3 MLP layers followed by **softmax** is used. The predicted probabilities of the extrusion sequence parameters $\tilde{\mathbf{e}}^* \in [0, 1]^{11 \times d_q \times L_e}$ are compared to the ground truth one-hot-encoded parameters in \mathbf{e}^* using a cross-entropy loss, \mathcal{L}_e .

Loop Decoder: Similarly to the extrusion decoder, the loop decoder Φ_ρ predicts the quantized parameters of the loop sequence $\rho^* \in \llbracket 0..d_q \rrbracket^{6 \times n_p \times L_\rho}$ as explained in Section 3. Nevertheless, 4 layers of multi-head transformer blocks are employed instead of simple MLP layers. This is due to the sequential nature of loop decoding in contrast to extrusions. Note that a similar strategy

as loop-extrusion decoder is opted for the transformer block of loop decoder. The first block performs self-attention on learned constant embedding of loops $\mathbf{F}_c^\rho \in \mathbb{R}^{d_z \times n_p L_\rho}$ and the result cross-attends to loop embedding \mathbf{F}_ρ as in Eq.(2). The same self and cross attention operations in Eq.(3) are conducted in subsequent blocks to yield a final representation at the last block $\mathbf{F}_\rho^b \in \mathbb{R}^{d_z \times n_p L_\rho}$. A linear layer followed by `softmax` is used to obtain predicted probabilities for the loop sequence parameters $\tilde{\rho}^* \in [0, 1]^{6 \times d_q \times n_p L_\rho}$ which are compared to ground truth one-hot-encoded loop parameters of ρ^* using a cross-entropy loss, \mathcal{L}_ρ .

Loop Refiner: As in many transformer-based architectures [3, 37, 43, 47], the quantization of loop parameters helps to reduce the search space and facilitates the learning. However, it has been observed in our case that it can lead to accumulation of quantization approximation errors. To overcome this issue, unquantized ground truth loop parameters are leveraged. In particular, a loop refiner Φ_r composed of a 4 layer MLP is introduced. This refiner takes as input a concatenation of loop embedding \mathbf{F}_ρ and their corresponding predicted parameter probabilities $\tilde{\rho}^*$. It attempts to predict the offset $\hat{\mathbf{O}}_{ff} \in \mathbb{R}^{6 \times n_p L_\rho}$ between the predicted quantized loop parameters $\hat{\rho}^* \in [0..d_q]^{6 \times n_p L_\rho}$ and the unquantized ground truth loop parameters $\rho^* \in \mathbb{R}^{6 \times n_p L_\rho}$. An MSE loss, \mathcal{L}_r , is computed between the predicted offset $\hat{\mathbf{O}}_{ff}$ and the one given by $\mathbf{O}_{ff} = \rho^* - \hat{\rho}^*$, to supervise the refiner and the rest of the network. Once the offset is predicted, it is added to the predicted quantized loop parameters yielding unquantized predicted loop parameters as follows $\hat{\rho} = \hat{\rho}^* + \hat{\mathbf{O}}_{ff}$.

Total Loss: *TransCAD* is an end-to-end network with a training objective guided by the sum of the individual losses, $\mathcal{L}_{total} = \mathcal{L}_{\rho,e} + \mathcal{L}_\rho + \mathcal{L}_e + \mathcal{L}_r$.

5 Proposed Evaluation

In this section, we first outline the limitations of existing evaluation methods in CAD sequence. Then, our new proposed evaluation metric framework for assessing the performance of CAD sequence inference from point clouds is described.

DeepCAD [43] Evaluation for Feature-based Reverse Engineering: An evaluation framework for CAD sequence was originally introduced in [43] and later used in [27]. This framework includes both accuracy for assessing the fidelity of the predicted sequence and Chamfer Distance (*CD*) to measure the quality of the recovered 3D geometry. Accuracy is assessed using two metrics, specifically *Command Type Accuracy* (ACC_{cmd}) and *Parameter Accuracy* (ACC_{param}) defined by

$$ACC_{cmd} = \frac{1}{N_c} \sum_{i=1}^{N_c} \mathbb{I}[t_i = \hat{t}_i], \quad (4)$$

$$ACC_{param} = \frac{1}{K} \sum_{i=1}^{N_c} \sum_{j=1}^{|\hat{\mathbf{p}}_i|} \mathbb{I}[|\mathbf{p}_{i,j} - \hat{\mathbf{p}}_{i,j}| < \eta] \mathbb{I}[t_i = \hat{t}_i], \quad (5)$$

where t_i and \hat{t}_i are the ground truth and predicted command types (for commands representing primitives and extrusions), $\mathbf{p}_{i,j}$ and $\hat{\mathbf{p}}_{i,j}$ are ground truth and predicted command parameters, N_c denotes the total number of CAD commands and $\mathbb{I}[\cdot]$ is the indicator function. $K = \sum_{i=1}^{N_c} \mathbb{I}[t_i = \hat{t}_i] |\mathbf{p}_i|$ is the total number of parameters of the correctly recovered commands and η is a tolerance threshold. The 3D geometry is evaluated with Chamfer Distance (CD) computed by sampling 2000 points on the ground truth and predicted shapes.

Limitations: We identify the following limitations of the aforementioned evaluation. **(1)** The proposed ACC_{cmd} overlooks the possibility of over-prediction in the CAD sequence. As indicated in Eq.(4), the computation of this metric sums across the set of ground truth CAD commands N_c . A predicted sequence could erroneously include extra loop-extrusion operations and still achieve a full score, as exemplified on the left panel of Fig. 3. **(2)** The evaluation of ACC_{param} is conducted solely on the subset of K accurately identified commands, thus introducing a trade-off between ACC_{param} and ACC_{cmd} . This interdependence complicates the interpretation of results. **(3)** Assessment of parameter quality via ACC_{param} solely in terms of accuracy is failing to distinguish between the magnitudes of errors. A parameter inaccurately placed in an adjacent quantization bucket incurs the same penalty as one with a larger deviation, despite potentially minor implications on the CAD model’s final geometry. These limitations cannot be entirely mitigated by complementing CAD command accuracies with the chamfer distance (CD) metric. While CD is a valuable assessment of shape similarity, it does not address the core objective of reverse engineering: to accurately recover the designer’s original CAD sequence. Two CAD models might be close in terms of CD yet possess vastly different CAD construction steps (see right panel of Fig. 3).

Proposed Evaluation Framework: To overcome the identified challenges, we introduce the *mean Average Precision of CAD Sequence* (APCS), a novel evaluation metric tailored for feature-based reverse engineering. APCS adopts the concept of Average Precision (AP) commonly used in other tasks, to quantify the similarity between predicted and ground truth CAD sequences. We introduce the *CAD Sequence Similarity Score* (CSSS) that can be computed between predicted and ground truth CAD sequences as follows

$$CSSS(\hat{\mathbf{C}}, \mathbf{C}) = \frac{1}{2N_\delta} \sum_{j=1}^{N^\rho} \sum_{i=1}^{N_j^\rho} \left[S(\hat{\delta}_{j,i}, \delta_{j,i}) \cdot \mathbb{I}[t_{yp}(\hat{\delta}_{j,i}) = t_{yp}(\delta_{j,i})] \right] + \frac{1}{2N^e} \sum_{j=1}^{N^e} S(\hat{\mathbf{e}}_j, \mathbf{e}_j), \quad (6)$$

where $t_{yp}(\cdot)$ is a function that determines the type of each primitive δ (arc, line, *etc.*), and $S(\hat{\mathbf{p}}, \mathbf{p}) = e^{-k\|\hat{\mathbf{p}}-\mathbf{p}\|}$ is a scoring function with $S(\hat{\mathbf{p}}, \mathbf{p}) \in [0, 1]$ with 1 assigned when predicted parameterization is identical to the ground truth. We define $N_j^\rho = \max(|\rho_j|, |\hat{\rho}_j|)$ where $|\cdot|$ denotes set cardinality, $N^\rho = \max(L_\rho, \hat{L}_\rho)$ where L_ρ is the number of predicted primitives and \hat{L}_ρ is the number of ground truth primitives for loop ρ_j and $N_\delta = \sum_{j=1}^{N^\rho} \max(|\rho_j|, |\hat{\rho}_j|)$. Finally, $N^e = \max(L_e, \hat{L}_e)$ where L_e is the number of predicted extrusions and

\hat{L}_e is the number of ground truth extrusions. The proposed CSSS metric evaluates both the operation type and parameter prediction. It assigns a score of 0 to loops with incorrectly predicted types, which gradually increases to 1 as parameter prediction improves. Assessment is conducted on the unquantized parameter space and calculates the score based on the maximum count of either predicted or ground truth primitives. This approach ensures that both over and under predicted sequences are penalized equally. We aggregate CSSS scores across various thresholds to derive the *mean Average Precision of CAD Sequences* (APCS). Furthermore, the median *CD* is used to measure shape similarity as in [43] with




Fig. 3: Two examples outlining the limitations of existing evaluation metrics. Left panel: the ground truth sequence (one sketch-extrusion) is a correct subset of the predicted sequence (three sketch-extrusions). The DeepCAD [43] metrics result in an accuracy of 1 for both commands and parameters. On the other hand, our proposed metric takes into account the over predicted sequence elements and the APCS is 0.031. The right panel showcases the limitations of the *CD* as a similarity measure. While the ground truth and predicted shapes are both composed of three extruded circle sketches, they are different in shape. However, the *CD* between the two shapes falls within the uncertainty range of ± 0.3 . Note that the uncertainty in the *CD* measurement is estimated by taking the average *CD* between all the test samples and themselves.

the difference that it is evaluated on 4096 points instead of 2000 in order to decrease the uncertainty in the *CD* measurement. All the reported *CD* measurements in this work are multiplied by 10^3 . Moreover, the ratio of predictions that cannot be reconstructed using [30] is reported as the invalidity ratio, IR.

6 Experiments

In this section, the experimental setup is first presented. Then, qualitative and quantitative results are analyzed. Afterwards, the components of *TransCAD* are ablated. Finally, the limitations of our model are outlined.

6.1 Experimental Setup

Dataset: For training and evaluation, the DeepCAD dataset [43] is used. The sketch extrusion sequences of the CAD models are processed in quantized (8 bits)

and unquantized space. The size of the train, validation and test sets are 140 294, 7 773, and 7 036 CAD models, respectively. Moreover, cross-dataset evaluation is conducted on the Fusion360 dataset [41] that contains 6 794 samples.

Training Details: The network is trained for 100 epochs with a batch size of 72 and an Adam optimizer is employed with a learning rate of 0.001 and a linear warm-up period of 2 000 steps as in [43]. The training is conducted on an NVIDIA RTX A6000 GPU. The input point clouds are extracted using [30] and are made of $n = 4096$ points. The dimension of the point features d_p is set to 16 and of loop-extrusion features d_z to 256.

Baselines: In order to evaluate the performance of *TransCAD*, two state-of-the-art methods, MultiCAD [27] and DeepCAD [43], and a retrieval baseline are used. As the code for MultiCAD [27] is not available, we report the results from the original paper. DeepCAD [43] is retrained with the same parameters and procedure as outlined in the original paper. One of the known limitations of the DeepCAD dataset is that it contains duplicate models [47]. While the works in [47] proposed a method to remove duplicate models that contain exactly the same CAD sequence, we find that this method does not remove all the duplicates as some models can have the same geometry but are constructed through a slightly different sequence of sketch extrusion operations (see supplementary material for more details). In order to address this issue, we propose a retrieval baseline. The retrieval baseline uses the point cloud encoder from a trained DeepCAD [43] to identify the closest latent vector from the train set for each test sample. As a result, the solution is always a train set CAD sequence.

6.2 Experimental Results

Qualitative Results: Fig. 4 shows some qualitative results for the retrieval baseline, DeepCAD [43] and *TransCAD* (Ours) on both the DeepCAD [43] and Fusion360 [41] datasets. As mentioned in Section 6.1, the DeepCAD dataset contains many duplicates, not just in terms of CAD sequence but also in terms of geometrical shape. As a result, the retrieval baseline is able to identify accurately duplicates (most right column of the DeepCAD dataset panel) and in other cases the baseline manages to retrieve shapes with similar geometry as the ground truth CAD model. On the other hand, it can be noticed that DeepCAD [43] can even fail at retrieving duplicates. *TransCAD* is able to predict models that are similar in shape and also in terms of loop-extrusion sequence, even though it sometimes fails to predict the parameters accurately (second and fifth columns of DeepCAD dataset and fourth column of Fusion360 dataset).

Quantitative Results: The trends observed from the qualitative results are further supported by the quantitative results presented in Table 1. The APCS, on both DeepCAD [43] and Fusion360 [41] datasets, show that *TransCAD* is the most capable model at predicting correct CAD sequences. However, it can be noted that the retrieval baseline obtains the lowest CD by a small margin on the DeepCAD dataset and by a more significant margin on the Fusion360. One of




Fig. 4: Qualitative results on the DeepCAD [43] dataset (left) and the cross-dataset Fusion360 [41] experiment (right).

DeepCAD					Fusion360				
	Model	APCS \uparrow	$CD\downarrow$	IR \downarrow		Model	APCS \uparrow	$CD\downarrow$	IR \downarrow
DeepCAD	<i>Retrieval</i>	0.629	2.8	0	Fusion360	<i>Retrieval</i>	0.304	60.0	0
	<i>MultiCAD</i> [27]	-	8.1	0.115		<i>MultiCAD</i> [27]	-	42.2	0.165
	<i>DeepCAD</i> [43]	0.604	19.2	0.038		<i>DeepCAD</i> [43]	0.360	104.2	0.017
	<i>Ours</i>	0.732	4.5	0.011		<i>Ours</i>	0.365	33.3	0.024

Table 1: Quantitative results on the DeepCAD [43] dataset and cross-dataset experiment on Fusion360 [41]. The APCS results show that *TransCAD* (*Ours*) is able to recover CAD sequence most accurately.

the reasons is that this baseline always outputs a CAD model that is of roughly similar shape as the input even if the retrieved CAD sequence can vary from the ground truth. To further analyse the results, the variations of the APCS and CD w.r.t. model complexity on the DeepCAD dataset are displayed in Fig. 5. We define the model complexity as the lowest possible CD of a test point cloud sample with respect to the train samples. In other words, the model complexity quantifies the amount by which a test sample is out of distribution from the train set in terms of shape. While *TransCAD* consistently outperforms on average all other baselines in terms of APCS for all model complexities, DeepCAD [43] can only perform better than the retrieval baseline for the more complex models. This shows that DeepCAD [43] often fails at retrieving the CAD sequence for simple models. In terms of CD , it can be observed that *TransCAD* and the retrieval baseline have similar performance. *TransCAD* can predict a CAD sequence that is closer to the ground truth one but the predicted overall shape can vary from the ground truth for more difficult samples.

Ablation Study: In this section, the different components of the proposed network architecture are ablated. Table 2 shows the results for *Ours w/o hier.*, in which the learning is done without both the loop-extrusion decoder and the refining network, *Ours w/o refining* where the refining component is ablated and *Ours*. It can be noted that each component leads to an improvement in




Fig. 5: Left: Plot of the variation of the mean APCS w.r.t. model complexity. Right: Plot of the variation of the median CD w.r.t. model complexity. For both graphs the number of models in each bin, represented by the horizontal bounded lines, correspond to approximately the same number of test samples from the DeepCAD dataset [43].

Model	APCS \uparrow	CD \downarrow	IR \downarrow
<i>Ours w/o hier.</i>	0.687	7.0	0.016
<i>Ours w/o refining</i>	0.708	4.8	0.018
<i>Ours</i>	0.732	4.5	0.011

Table 2: Ablation results demonstrating the relevance of the hierarchical learning strategy and the loop refiner.

all the metrics. It is worth noting that the $F1$ score on the loop-extrusion type prediction introduced for the hierarchical learning of *TransCAD* is 0.79. This implies that on most cases both loop and extrusion decoders receive embedding of the correct type. Moreover, while the refining network is only applied on the loop parameters, we observe that the component of the APCS for the extrusion parameters are also higher when the network is trained with the refining component. This suggests that the refining network is able to provide a useful signal for guiding the learning process. More details are in the supplementary material.

Input point cloud perturbation: Reverse engineering is a real-world practical problem. The results in the previous section are obtained from sampling points from the B-Rep representations of the CAD models. While modern 3D sensors can reconstruct the mesh of models with high resolution, they still suffer from some artifacts such as noise and small missing parts. In order to evaluate the performance of our network in such realistic conditions, we run experiments in two scenarios, one in which noise is added and one in which small holes are created on the point cloud. In order to simulate realistic noise, Perlin noise [31] is added to the mesh from which the point coordinates and normals are extracted. More details about the noise and hole generations can be found in the supplementary materials. Table 3 shows that *TransCAD* is more robust to such perturbations than other methods. As the noise also adds a disturbance to the direction of the input point normals, it leads to a larger drop in performance.

	Model	APCS \uparrow	CD \downarrow	IR \downarrow		Model	APCS \uparrow	CD \downarrow	IR \downarrow
Noise	<i>Retrieval</i>	0.561	5.9	0	Holes	<i>Retrieval</i>	0.607	4.1	0
	<i>DeepCAD</i> [43]	0.550	31.1	0.047		<i>DeepCAD</i> [43]	0.574	26.7	0.039
	<i>Ours</i>	0.604	18.1	0.010		<i>Ours</i>	0.732	4.4	0.012

Table 3: Results on the DeepCAD [43] dataset when the input point cloud is perturbed either with Perlin noise (left) or by creating holes (right).

Failure Cases and Limitations: In this section, we describe the reasons that lead *TransCAD* to predict invalid CAD models as measured by IR. Among the predictions of *TransCAD* on DeepCAD test set, only one contains a loop parametrization that results in an invalid CAD sequence. This shows that the proposed representation of the loop sequence leads to syntactically correct loops on practically all cases. However, in some cases the representation of the CAD sequence can be syntactically valid, yet it is not possible to reconstruct a B-Rep from it. For 75 test samples, the loop-extrusion decoder fails to predict an extrusion token, this implies that the predicted model is therefore an infinitely thin sketch and not a 3D model as expected. The rest of the invalid models are mostly due to a loop being made of a single line within a model, which cannot be extruded into a valid 3D shape within our context. Finally, Fig. 6 shows examples for which the predicted sequences do not lead to a shape that is close to the ground truth one. In these examples, the input CAD models contain a large number of small features that *TransCAD* is unable to capture.




Fig. 6: Examples for which *TransCAD* fails at recovering a shape close to the ground truth one.

7 Conclusion

In conclusion, we propose *TransCAD* an end-to-end transformer-based neural network that learns to recover the CAD sequence from a given point cloud. Two of the main features of *TransCAD* are a hierarchical structure that enables the learning of a high-level loop-extrusion sequence and a loop refiner that aims at correcting errors in loop parameter predictions. We also propose a primitive representation in which each primitive is described by the same number of parameters. We identify the limitations of current metrics in the emerging domain of 3D reverse engineering and propose a new metric, the APCS that leads to a fair comparison of parametric CAD sequences. Thorough experiments show that *TransCAD* achieves state-of-the-art results in different realistic scenarios.

Acknowledgement: The present project is supported by the National Research Fund, Luxembourg under the BRIDGES2021/IS/16849599/FREE-3D and IF/17052459/CASCADES projects, and by Artec 3D.

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: ICML. pp. 40–49. PMLR (2018)
2. Cai, R., Yang, G., Averbuch-Elor, H., Hao, Z., Belongie, S., Snavely, N., Hariharan, B.: Learning gradient fields for shape generation. In: ECCV (2020)
3. Carlier, A., Danelljan, M., Alahi, A., Timofte, R.: Deepsvg: A hierarchical generative network for vector graphics animation. *NeurIPS* **33**, 16351–16361 (2020)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
5. Chang, H., Zhang, H., Jiang, L., Liu, C., Freeman, W.T.: Maskgit: Masked generative image transformer. In: CVPR. pp. 11315–11325 (2022)
6. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: CVPR. pp. 5939–5948 (2019)
7. Cherenkova, K., Dupont, E., Kacem, A., Arzhannikov, I., Gusev, G., Aouada, D.: Sepicnet: Sharp edges recovery by parametric inference of curves in 3d shapes. In: CVPRW. pp. 2726–2734 (2023)
8. Dupont, E., Cherenkova, K., Kacem, A., Ali, S.A., Aryhannikov, I., Gusev, G., Aouada, D.: Cadops-net: Jointly learning cad operation types and steps from boundary-representations. *3DV* (2022)
9. Friedrich, M., Fayolle, P.A., Gabor, T., Linnhoff-Popien, C.: Optimizing evolutionary csg tree extraction. In: Proceedings of the Genetic and Evolutionary Computation Conference. pp. 1183–1191 (2019)
10. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial networks. *NeurIPS* (2014)
11. Guo, H., Liu, S., Pan, H., Liu, Y., Tong, X., Guo, B.: Complexgen: Cad reconstruction by b-rep chain complex generation. *ACM TOG* **41**(4), 1–18 (2022)
12. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *NeurIPS* (2020)
13. Jayaraman, P.K., Lambourne, J.G., Desai, N., Willis, K.D., Sanghi, A., Morris, N.J.: Solidgen: An autoregressive model for direct b-rep synthesis. *Transaction in Machine Learning Research* (2023)
14. Jayaraman, P.K., Sanghi, A., Lambourne, J.G., Willis, K.D., Davies, T., Shayani, H., Morris, N.: Uv-net: Learning from boundary representations. In: CVPR. pp. 11703–11712 (2021)
15. Kania, K., Zieba, M., Kajdanowicz, T.: Ucs-g-net-unsupervised discovering of constructive solid geometry tree. *NeurIPS* (2020)
16. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020)
17. Khan, M.S., Dupont, E., Ali, S.A., Cherenkova, K., Kacem, A., Aouada, D.: Cad-signet: Cad language inference from point clouds using layer-wise sketch instance guided attention. In: CVPR. pp. 4713–4722 (2024)
18. Koch, S., Matveev, A., Jiang, Z., Williams, F., Artemov, A., Burnaev, E., Alexa, M., Zorin, D., Panozzo, D.: Abc: A big cad model dataset for geometric deep learning. In: CVPR. pp. 9601–9611 (2019)

19. Lambourne, J.G., Willis, K.D., Jayaraman, P.K., Sanghi, A., Meltzer, P., Shayani, H.: Brepnet: A topological message passing system for solid models. In: CVPR. pp. 12773–12782 (2021)
20. Lambourne, J.G., Willis, K., Jayaraman, P.K., Zhang, L., Sanghi, A., Malekshan, K.R.: Reconstructing editable prismatic cad from rounded voxel models. In: SIGGRAPH Asia. pp. 1–9 (2022)
21. Li, J., Xu, K., Chaudhuri, S., Yumer, E., Zhang, H., Guibas, L.: Grass: Generative recursive autoencoders for shape structures. ACM TOG (2017)
22. Li, L., Sung, M., Dubrovina, A., Yi, L., Guibas, L.J.: Supervised fitting of geometric primitives to 3d point clouds. In: CVPR. pp. 2652–2660 (2019)
23. Li, P., Guo, J., Zhang, X., Yan, D.M.: Secad-net: Self-supervised cad reconstruction by learning sketch-extrude operations. In: CVPR. pp. 16816–16826 (2023)
24. Lim, J.J., Pirsivash, H., Torralba, A.: Parsing ikea objects: Fine pose estimation. In: CVPR. pp. 2992–2999 (2013)
25. Liu, Y., D’Aronco, S., Schindler, K., Wegner, J.D.: Pc2wf: 3d wireframe reconstruction from raw point clouds. ICLR (2021)
26. Ma, W., Chen, S., Lou, Y., Li, X., Zhou, X.: Draw step by step: Reconstructing cad construction sequences from point clouds via multimodal diffusion. In: CVPR. pp. 27154–27163 (2024)
27. Ma, W., Xu, M., Li, X., Zhou, X.: Multicad: Contrastive representation learning for multi-modal 3d computer-aided design models. In: CIKM. Association for Computing Machinery, New York, NY, USA (2023)
28. Mallis, D., Aziz, A.S., Dupont, E., Cherenkova, K., Karadeniz, A.S., Khan, M.S., Kacem, A., Gusev, G., Aouada, D.: Sharp challenge 2023: Solving cad history and parameters recovery from point clouds and 3d scans. overview, datasets, metrics, and baselines. In: CVPRW (2023)
29. Nash, C., Ganin, Y., Eslami, S.A., Battaglia, P.: Polygen: An autoregressive generative model of 3d meshes. In: ICML. pp. 7220–7229. PMLR (2020)
30. Paviot, T.: Pythonocc (2017), <https://dev.opencascade.org/project/pythonocc>, accessed: March 7, 2024
31. Perlin, K.: An image synthesizer. In: SIGGRAPH. p. 287–296. SIGGRAPH ’85, Association for Computing Machinery, New York, NY, USA (1985)
32. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. NeurIPS **30** (2017)
33. Ren, D., Zheng, J., Cai, J., Li, J., Zhang, J.: Extrudenet: Unsupervised inverse sketch-and-extrude for shape parsing. In: ECCV. pp. 482–498. Springer (2022)
34. Sharma, G., Liu, D., Maji, S., Kalogerakis, E., Chaudhuri, S., Měch, R.: Parsenet: A parametric surface fitting network for 3d point clouds. In: ECCV. pp. 261–276. Springer (2020)
35. Team, O.D.: Open3d documentation: `open3d.geometry.estimate_normals` (2022), https://www.open3d.org/docs/0.7.0/python_api/open3d.geometry.estimate_normals.html, accessed: March 6, 2024
36. Uy, M.A., Chang, Y.Y., Sung, M., Goel, P., Lambourne, J.G., Birdal, T., Guibas, L.J.: Point2cyl: Reverse engineering 3d objects from point clouds to extrusion cylinders. In: CVPR. pp. 11850–11860 (2022)
37. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS **30** (2017)
38. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: ECCV. pp. 52–67 (2018)

39. Wang, X., Xu, Y., Xu, K., Tagliasacchi, A., Zhou, B., Mahdavi-Amiri, A., Zhang, H.: Pie-net: Parametric inference of point cloud edges. *NeurIPS* **33**, 20167–20178 (2020)
40. Wijmans, E.: Pointnet++ pytorch (2018), https://github.com/erikwijmans/Pointnet2_PyTorch, accessed: March 7, 2024
41. Willis, K.D., Pu, Y., Luo, J., Chu, H., Du, T., Lambourne, J.G., Solar-Lezama, A., Matusik, W.: Fusion 360 gallery: A dataset and environment for programmatic cad construction from human design sequences. *ACM TOG* **40**(4), 1–24 (2021)
42. Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. *NeurIPS* **29** (2016)
43. Wu, R., Xiao, C., Zheng, C.: Deepcad: A deep generative network for computer-aided design models. In: *CVPR*. pp. 6772–6782 (2021)
44. Wu, R., Zhuang, Y., Xu, K., Zhang, H., Chen, B.: Pq-net: A generative part seq2seq network for 3d shapes. In: *CVPR* (2020)
45. Xu, X., Jayaraman, P.K., Lambourne, J.G., Willis, K.D., Furukawa, Y.: Hierarchical neural coding for controllable cad model generation. *ICML* (2023)
46. Xu, X., Lambourne, J.G., Jayaraman, P.K., Wang, Z., Willis, K.D., Furukawa, Y.: Brepgen: A b-rep generative diffusion model with structured latent geometry. *ACM SIGGRAPH* (2024)
47. Xu, X., Willis, K.D., Lambourne, J.G., Cheng, C.Y., Jayaraman, P.K., Furukawa, Y.: Skexgen: Autoregressive generation of cad construction sequences with disentangled codebooks. In: *ICML*. pp. 24698–24724. PMLR (2022)
48. Xu, X., Peng, W., Cheng, C.Y., Willis, K.D., Ritchie, D.: Inferring cad modeling sequences using zone graphs. In: *CVPR*. pp. 6062–6070 (2021)
49. Yang, G., Huang, X., Hao, Z., Liu, M.Y., Belongie, S., Hariharan, B.: Pointflow: 3d point cloud generation with continuous normalizing flows. In: *CVPR*. pp. 4541–4550 (2019)
50. Yu, F., Chen, Q., Tanveer, M., Mahdavi Amiri, A., Zhang, H.: D2csg: Unsupervised learning of compact csg trees with dual complements and dropouts. *Advances in Neural Information Processing Systems* **36** (2024)
51. Yu, L., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Ec-net: an edge-aware point set consolidation network. In: *ECCV*. pp. 386–402 (2018)
52. Zhu, X., Du, D., Chen, W., Zhao, Z., Nie, Y., Han, X.: Nerve: Neural volumetric edges for parametric curve extraction from point cloud. In: *CVPR*. pp. 13601–13610 (2023)

Supplementary Materials

TransCAD: A Hierarchical Transformer for CAD Sequence Inference from Point Clouds

A Formulation

Fig. 1 shows an example of the construction process of a CAD model as well as the corresponding loop-extrusion sequences. The loop-extrusion sequence is made of three types of tokens (loop, extrusion and end of sequence). The CAD sequence combines the high level loop-extrusion sequence tokens and their corresponding parameters. Note for the loop parameters, the coordinates of the start point of a primitive is always the same as the coordinates of the end point of the previous primitive as in [43].




Fig. 1: Example of a CAD model and its CAD sequence in the formulation proposed. The top panels depicts the CAD construction process. The middle panel shows the corresponding high-level loop-extrusion sequence that is predicted as part of the hierarchical learning. The low-level parameters of each loop primitive and extrusion are displayed in the bottom panel.

B Duplicates in the DeepCAD Dataset

As mentioned in Section 6.1 of the main paper, one of the main limitations of the DeepCAD dataset [43] is that it contains many duplicates across the train and

test sets. While the works in [47] proposed a strategy to identify duplicates that have exactly the same CAD sequence (*i.e. sequence duplicates*), we observe that the dataset also contains models that have almost identical geometry but different CAD sequences (*i.e. geometrical duplicates*). Moreover some CAD models in the DeepCAD train set are almost identical to some models in the Fusion360 dataset [41], with often just the amount of extrusion varying slightly. Fig. 2 shows examples of the different types of duplicates from the DeepCAD [43] and Fusion360 [41] datasets. We define a geometrical duplicate as a test set CAD model for which it exists a CAD model in the train set with a chamfer distance less than the uncertainty in the chamfer distance measurement ($\pm 3 \times 10^{-4}$ when 4096 points are sampled). From this definition and using the train set of the DeepCAD dataset [43], we observe that about 14% of the DeepCAD test set is made of geometrical duplicates and about 12% in the Fusion360 test set. We observe that current datasets that contain both CAD models and their corresponding CAD sequences are limited by either by the number of samples (Fusion360 [41]) or by the lack of diversity they present (DeepCAD [43]).

C Implementation details

In this section, more details on the network architecture are provided.

Transformer decoders: The loop-extrusion decoder $\Phi_{\rho,e}$ and the loop decoder Φ_{ρ} are both transformer decoder with the same network architecture. They are both made of 4 layers, each made of 8 heads with a feed-forward dimension of 512. A dropout rate of 0.1 is used.

PointNet++: As mentioned in the main paper, the point encoder Φ_p is PointNet++ [32]. The implementation provided in [40] was used. The input dimension $d_p = 6$ corresponds to the point and normal coordinates. The parameters for the 4 layers are as follows: number of points (512, 256, 128, 16) with radius (0.1, 0.2, 0.4, 0.8) and number of samples (64, 64, 64, 32).

D APCS Metric

In this section, more details on the computation of the proposed metric APCS (see Section 5 of the main paper) are provided.

The *mean Average Precision of CAD Sequence*, APCS, is a score between 0 and 1 that evaluates how close two CAD sequences are by taking into account the types and parameters of loop primitives and the extrusion parameters. The types of primitives considered in this work are *line*, *arc* and *circle*. The parameters of the loop primitives correspond to 6 point coordinates in a normalized 2D space. An extrusion is described using 11 parameters as in [43] which can be grouped into 4 main categories: 1) Ext.: the type of extrusion and the amount of extrusion; 2) Origin: the origin of the 2D sketch in 3D space; 3) Orientation: the sketch plane orientation; and 4) Size: the size of the sketch in 3D space. Note the latter




Fig. 2: Examples of duplicate CAD models from the DeepCAD [43] and Fusion360 [41] datasets. On the top left panel, CAD models from the DeepCAD train set with geometrical duplicates in the test set are shown. Similarly, the right panel presents geometrical duplicates present in the Fusion360 [41] dataset. CAD models with identical CAD sequences, *i.e.* sequence duplicates, are displayed in the bottom left panel.

3 categories describe the projection and scaling of the normalized 2D sketch in 3D.

E Results

In this section, further quantitative and qualitative results are presented.

E.1 Quantitative Results

The APCS scores reported in the main paper are computed per model and then averaged over all the models of the test set. It is also possible to compute the average of each individual component of the APCS scores over the whole test set. Such results are discussed in the following paragraphs for the DeepCAD dataset [43], Fusion360 dataset [41] and the ablation study.

Model	APCS \uparrow						
	Line	Arc	Circle	Ext.	Origin	Orientation	Size
<i>Retrieval</i>	0.584	0.280	0.655	0.716	0.666	0.819	0.691
<i>DeepCAD</i> [43]	0.654	0.246	0.587	0.872	0.825	0.928	0.848
<i>Ours</i>	0.665	0.709	0.683	0.818	0.768	0.879	0.778

Table 1: Results of the different APCS components on the DeepCAD dataset [43].

Model	$ACC_{cmd}\uparrow$	$ACC_{param}\uparrow$
<i>DeepCAD</i> [43]	0.821	0.693
<i>Ours</i>	0.838	0.741

Table 2: Results using the metrics described in DeepCAD on the DeepCAD dataset [43].

Model	APCS \uparrow						
	Line	Arc	Circle	Ext.	Origin	Orientation	Size
<i>Retrieval</i>	0.214	0.045	0.546	0.453	0.294	0.509	0.451
<i>DeepCAD</i> [43]	0.368	0.144	0.639	0.844	0.774	0.866	0.859
<i>Ours</i>	0.515	0.383	0.622	0.749	0.693	0.836	0.782

Table 3: Results of the different APCS components on the Fusion dataset [41].

DeepCAD Dataset: Table 1 shows the APCS scores on the DeepCAD dataset [43] for each of the components averaged over the test set. It can be observed that *TransCAD* (*Ours*) obtains a significantly better score for the arc primitive and also to some extent for the circle primitive compared to the retrieval baseline and DeepCAD [43]. However, the scores corresponding to the placement of the 2D sketch in 3D (Origin, Orientation and Size) for *TransCAD* are slightly lower than for DeepCAD [43].

The APCS score provides a more complete evaluation of the predicted CAD sequences than the command accuracy (ACC_{cmd}) and parameter accuracy (ACC_{param}) used in the works [43]. Nevertheless, we provide the results for *TransCAD* and DeepCAD [43] against those metrics in Table 2 for the DeepCAD dataset.

Fusion360 Dataset: The trends previously described on the DeepCAD dataset [43] can also be observed on the cross-dataset evaluation using the Fusion360 dataset [41] (see Table 3). Furthermore, the APCS score corresponding to the line primitive is significantly higher for *TransCAD* than for the other two baseline models. Fig. 3 shows the variation of the APCS and CD w.r.t. to model complexity for the Fusion360 dataset [41]. It can be observed that while the performance on the APCS metric for DeepCAD [43] and *TransCAD* are relatively close, *TransCAD* achieves lowest CD for all model complexities compared to DeepCAD [43] and the retrieval baseline.




Fig. 3: Left: Plot of the variation of the mean APCS w.r.t. model complexity. Right: Plot of the variation of the median CD w.r.t. model complexity. For both graphs the number of models in each bin, represented by the horizontal bounded lines, correspond to approximately the same number of test samples from the Fusion360 dataset [41].

Model	APCS \uparrow						
	Line	Arc	Circle	Ext.	Origin	Orientation	Size
<i>Ours w/o hier.</i>	0.569	0.473	0.529	0.749	0.640	0.843	0.670
<i>Ours w/o refining</i>	0.655	0.625	0.673	0.755	0.679	0.838	0.703
<i>Ours</i>	0.665	0.709	0.683	0.818	0.768	0.879	0.778

Table 4: Results of the different APCS components on the DeepCAD dataset [43] for the ablation study. *Ours w/o hier.* corresponds the proposed model without hierarchical learning, and Loop Refiner and *Ours w/o refining* to the proposed model without the Loop Refiner.

Complex shape performance: Figure 4 show the APCS w.r.t. the sequence length for *TransCAD* and DeepCAD [43]. Similarly, Figure 5 presents the variation of CD w.r.t. the sequence length. The size of the data points is proportional to the number of models it represents. The performance decreases for both models as the length of the CAD sequence increases. However, *TransCAD* consistently outperforms DeepCAD.

Ablation Study: Table 4 shows the APCS component scores corresponding to the ablation study presented in Section 6.2 of the main paper. The most striking result is that the Loop Refiner leads to improved performance in all the extrusion components, even though the Loop Refiner only acts directly on the loop parameters. The Loop Refiner is a component of the end to end pipeline of *TransCAD*. As a result, the backpropagation of the Loop Refiner loss \mathcal{L}_ρ acts on all the parameters of the network and can therefore impact the predictions at all levels.

The Loop-Extrusion module classifies the features $\mathbf{F}_{\rho,e}$ as *loop*, *extrusion* or *end of sequence* type. These predictions are used to route the features $\mathbf{F}_{\rho,e}$ to either the loop decoder Φ_ρ or extrusion decoder Φ_e to obtain their parameters.




Fig. 4: APCS results as a function of the ground truth CAD sequence length expressed in the DeepCAD format on the DeepCAD dataset [43]. The size of the points is proportional to the number of CAD models with the corresponding CAD sequence length.

To demonstrate the impact of the Loop-Extrusion type classification, we conduct the following experiment: the ground truth Loop-Extrusion type labels are used instead of the predicted ones at testing time on the DeepCAD dataset [43]. In this scenario, the APCS metric evaluating the CAD sequence increases from 0.732 to 0.790. The *IR* also improves and decreases to nearly 0% (only 2 invalid models). Notably, there is no significant change in the *CD*. As a result, *TransCAD* is robust to moderate classification errors in the Loop-Extrusion prediction w.r.t. the final reconstruction. However, these errors might impact the performance of the predicted CAD sequence w.r.t. the ground truth. This suggests that the Loop-Extrusion classification errors might result in alternative yet plausible design paths.

E.2 Qualitative Results

In this section, further qualitative results are presented by first comparing *TransCAD* to the retrieval baseline and then to DeepCAD [43].

Retrieval Baseline Comparison: Qualitative results comparing *TransCAD* to the the retrieval baseline can be found in Fig. 6. It can be observed that *TransCAD* can perform well on duplicate models and more importantly that it performs better than the retrieval baseline on non-duplicate models. *TransCAD* can identify most of the components of an unseen CAD model but sometimes




Fig. 5: CD results as a function of the ground truth CAD sequence length expressed in the DeepCAD format on the DeepCAD dataset [43]. The size of the points is proportional to the number of CAD models with the corresponding CAD sequence length.

fails to place those parts in the right location as suggested by the APCS scores presented in the previous section.

DeepCAD Comparison: Qualitative results showing the comparison between *TransCAD* and DeepCAD [43] can be found in Fig. 7. The results on *simple* and *difficult* models show that *TransCAD* is able to outperform DeepCAD [43] on most occasions.

F Comparison with ComplexGen

The output of *TransCAD* is a CAD sequence that allows full integration and editability in CAD software at both the final shape and intermediate design process levels. On the other hand, the output of ComplexGen [11] is drastically different since it is only the final shape as a set of corners, curves, and patches with topology constraints. ComplexGen cannot predict the intermediate design steps, as shown in Fig. 1 in the paper. To the best of our knowledge, DeepCAD [43] and MultiCAD [27] are the only works that predict CAD sequences from point clouds, with only DeepCAD providing source codes for extensive comparisons. Nevertheless, we ran a pretrained ComplexGen model on the DeepCAD test set and recorded a CD of 2.3 (compared to 4.5 for *TransCAD*). Note that as ComplexGen does not predict the CAD sequence, the APCS metric cannot be computed. On the other hand, ComplexGen metrics are tailored to their output




Fig. 6: Qualitative results showing the performance of *TransCAD* against the retrieval baseline on both duplicate CAD models (top panel) and non-duplicate models (bottom panel).

and designed to evaluate the predicted patches of the final shape and topological validity. To conclude, other reverse engineering baselines (*e.g.* ComplexGen [11]) might achieve a better final reconstruction than *TransCAD*, but could not be used seamlessly for reverse engineering (*i.e.* intermediate design level editability).

G Point Cloud Perturbation

In this section, further details on the implementation of the perturbations applied to the input point cloud \mathbf{P} described in Section 6.2 of the main paper are presented.

G.1 Perlin Noise Implementation

To simulate the noise created when an object is scanned using a 3D sensor, a Perlin noise [31] is added to the point cloud. The perlin noise is created using the following strategy. Starting from the mesh representation of the original CAD models, the faces are divided to ensure that the mesh contains a dense number of vertices. Then, a 3D Perlin noise is computed for each vertex using 64 octaves with a minimum and maximum magnitude of -0.001 and 0.001 , respectively. Finally, the normals are recomputed from the mesh and points are sampled. A visual example of a perturbed mesh can be found in Fig. 8.

G.2 Holes Implementation

Holes in the input point clouds are created using the following strategy. Firstly, for each point cloud the number holes is selected from a uniform distribution ranging from 1 to 10 included. Then, the ratio of points to be removed for each hole is chosen by sampling a normal distribution with mean 0.03 and standard deviation 0.015. Finally, for each hole a point is chosen at random and the corresponding number of nearest neighbors points are removed. The nearest neighbors are identified using a geodesic distance computed on the mesh surface. We ensure that the remaining number of points is at least $n = 4096$, which corresponds to the number of points used as input. Fig. 9 shows some examples of point clouds on which holes have been created.

G.3 Qualitative Results

Qualitative results for both the hole and noise input perturbations can be found in Fig. 10. Those results complement the quantitative results found in Section 6.2 of the main paper.

G.4 Further Quantitative Results

In this section, further quantitative results for different amounts of point cloud perturbation are presented. The maximum magnitude of the Perlin noise on the input point cloud is increased compared to the one reported in the main paper (now referred to as *Original Noise*). The results presented in Table 5 show that both *TransCAD* and DeepCAD are sensitive to the amount of noise. We also perform 2-epoch finetuning of both models on *Original Noise* training data and report the results (*Finetune*, last row of Table 5). Notably, *TransCAD* almost recovers its performance with finetuning.

Furthermore, we conduct an experiment to evaluate the effect of the input cloud sparsity on the performance. Using *TransCAD* and DeepCAD [43] trained with input point clouds of size 4096 points, predictions for the test set with decreasing input point cloud sizes are generated. The results demonstrate that *TransCAD* is more robust than DeepCAD [43] w.r.t. input sparsity (see Table 6).

Noise increase	<i>TransCAD</i>		DeepCAD [43]	
	<i>APCS</i> \uparrow	<i>CD</i> \downarrow	<i>APCS</i> \uparrow	<i>CD</i> \downarrow
<i>Original Noise</i>	0.604	18.1	0.550	31.1
+25%	0.570	21.1	0.541	31.6
+50%	0.532	28.8	0.528	35.6
<i>Finetune</i>	0.724	5.1	0.572	27.7

Table 5: Results on the *APCS* and *CD* metrics for different amount of noise added to the input point cloud on the DeepCAD dataset [43].

Input size (# Points)	<i>TransCAD</i>		DeepCAD[40]	
	<i>APCS</i> \uparrow	<i>CD</i> \downarrow	<i>APCS</i> \uparrow	<i>CD</i> \downarrow
4096	0.732	4.5	0.604	19.2
2048	0.729	4.9	0.577	24.8
512	0.705	8.1	0.526	38.5
256	0.673	12.7	0.443	70.3

Table 6: Results on the *APCS* and *CD* metrics for different size of input clouds on the DeepCAD dataset [43].

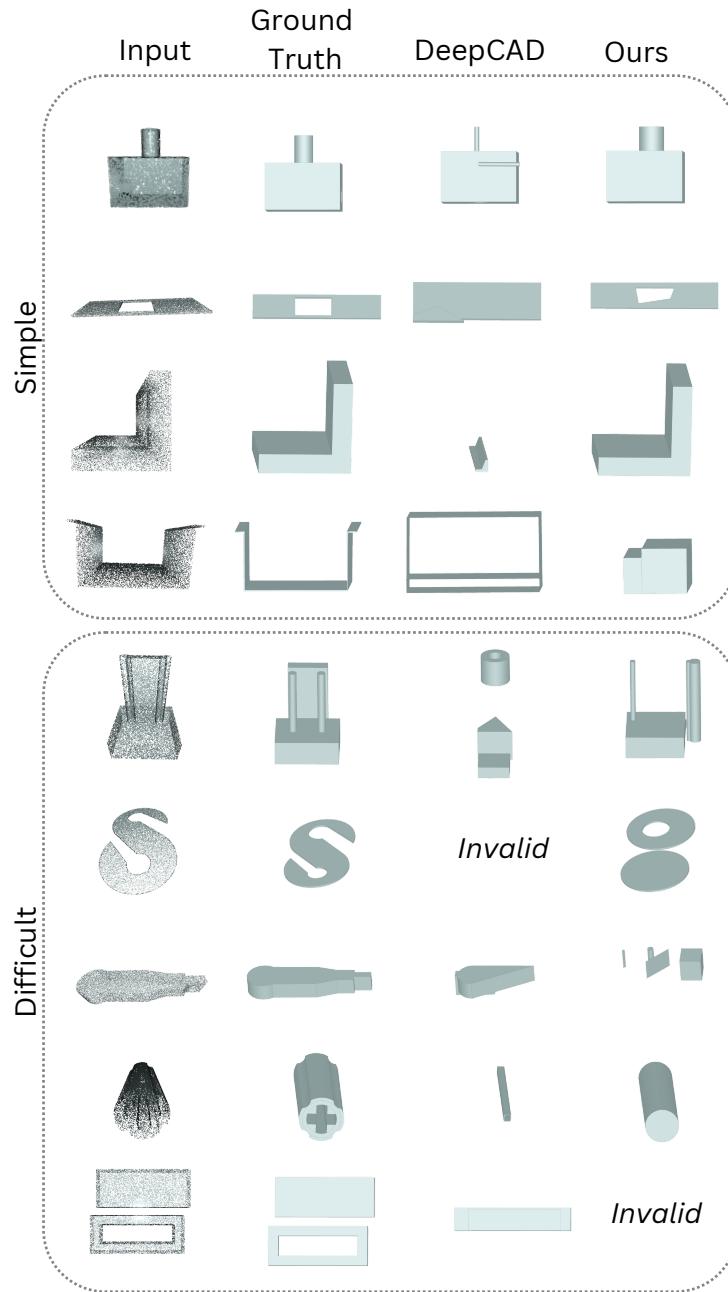


Fig. 7: Qualitative results showing the performance of *TransCAD* against the DeepCAD [43] on both simpler CAD models (top panel) and more complex models (bottom panel).




Fig. 8: Example of a CAD model with its mesh representation (left) and perturbed representation with Perlin noise (right). Zooming on the figure might be required to best view the effect of the Perlin noise applied.




Fig. 9: Examples of input point clouds in which holes have been created. The points highlighted in red represent the points that have been removed.




Fig. 10: Qualitative results on DeepCAD dataset [43] using perturbed input point clouds. The top panel shows results for which holes are created on the input point cloud and the bottom panel shows results for which Perlin noise was applied.