
Discovering Design Concepts for CAD Sketches

Anonymous Author(s)

Affiliation

Address

email

Abstract

Sketch design concepts are recurring patterns found in parametric CAD sketches. Though rarely explicitly formalized by the CAD designers, these concepts are implicitly used in design for modularity and regularity. In this paper, we propose a learning based approach that discovers the modular concepts by induction over raw sketches. We propose the dual implicit-explicit representation of concept structures that allows implicit detection and explicit generation, and the separation of structure generation and parameter instantiation for parameterized concept generation, to learn modular concepts by end-to-end training. We demonstrate the design concept learning on a large scale CAD sketch dataset and show its applications for design intent interpretation and auto-completion.

1 Introduction

Parametric CAD modeling is a standard paradigm for mechanical CAD design nowadays. In parametric modeling, CAD sketches are fundamental 2D shapes used for various 3D construction operations. As shown in Fig. 1, a CAD sketch is made of primitive geometric elements (e.g. lines, arcs, points) which are constrained by different relationships (e.g. coincident, parallel, tangent); the sketch graph of primitive elements and constraints captures design intents, and allows adaptation and reuse of designed parts by changing parameters and updating all related elements automatically [1]. Designers are therefore tasked with the meticulous design of such sketch graphs, so that the inherent high-level design intents are easy to interpret and disentangle. To this end, meta-structures (Fig. 1), which we call *sketch concepts* in this paper, capture repetitive design patterns and regulate the design process with more efficient intent construction and communication [9, 12]. Concretely, each sketch concept is a structure that encapsulates specific primitive elements and their compositional constraints, and the interactions of its internal elements with outside only go through the interface of the concept.

Is it possible to discover these modular concepts automatically from raw sketch graphs? In this paper, we cast this task as a program library induction problem by formulating a domain specific language (DSL) for sketch generation, where a sketch graph is formalized as a program, and sketch concepts are modular functions that abstract primitive elements and compose the program. Discovering sketch concepts thus becomes the induction of library functions from sketch programs. While previous works address the general library induction problem via expensive combinatorial search [19, 5, 7], we present a simple end-to-end deep learning solution for sketch concepts. Specifically, we bridge the implicit and explicit representations of sketch concepts, and separate concept structure generation from parameter instantiation, so that a powerful deep network can detect and generate sketch concepts, by training with the inductive objective of reconstructing sketch with modular concepts.

We conduct experiments on large-scale sketch datasets [16]. The learned sketch concepts show that they provide modular interpretation of design sketches. The network can also be trained on incomplete input sketches and learn to auto-complete them. Comparisons with state-of-the-art approaches that solve sketch graph generation through autoregressive models show that the modular sketch concepts learned by our approach enable more accurate and interpretable completion results.




Figure 1: **Concept learning from sketch graphs.** In black are the raw sketch and its constraint graph, with nodes showing primitives and edges depicting constraints. Colored are the restructured sketch and its modular constraint graph, where each module box represents a concept; primitives and constraint edges are colored according to the modular concepts.

2 Related work

Concept discovery for CAD sketch It is well acknowledged in the CAD design community that design intents are inherent to and implicitly encoded by the combinations of geometric primitives and constraints [12, 10]. However, there is generally no easy approach to discover the intents and make them explicit, albeit through manual design of meta-templates guided by expert knowledge [9, 10]. We propose an automatic approach to discover such intents, by formulating the intents as modular structures with self-contained references, and learning them through self-supervised inductive training with simple objectives on large raw sketch dataset. Therefore, we provide an automatic approach for discovering combinatorially complex structures through end-to-end neural network learning.

Generative models for CAD sketch A series of recent works [6, 23, 13, 17, 24] use autoregressive models [21] to generate CAD sketches and constraints modeled through pointer networks [22]. These works focus on learning from large datasets [16] to generate plausible layouts of geometric primitives and their constraints, which can then be fine-tuned with a constraint solver for more regular sketches. Different from these works, our aim is to discover modular structures (i.e. sketch concepts) from the concrete sketches. Therefore, our framework provides higher-level interpretation of raw sketches and more transparent auto-completion than these works (cf. Sec. 6).

Program library induction for CAD modeling Program library induction has been studied in the shape modeling domain [7]. General program synthesis assisted by deep learning is a research topic with increasing popularity [19, 3, 4, 18, 5]. The library induction task specifically involves combinatorial search, as has been handled by neural guided search [19, 5] or by pure stochastic sampling [7]. We instead present an end-to-end learning algorithm for sketch concept induction. In particular, based on key observations about sketch concepts, we present implicit-explicit dual representations of concept library functions, and separate the concept structure generation from parameter instantiation, to enable self-supervised training with induction objectives.

3 CAD sketch concept formulation

To capture the notion of sketch concepts precisely, we formulate a domain specific language (DSL) (syntax given in List 1, an exhaustive list of data types given in the supplementary). In the DSL, we first define the basic data types, including *length*, *angle*, *coordinate*, and the *reference* type, where a reference binds to another reference or a primitive for modeling the constraint relationships. Second, we define the \mathbb{L}^0 collection of primitive and constraint types as given in raw sketches. In particular, we regard the constraints as functions whose arguments are the references to bind with primitives, e.g. a coincident constraint $c = \lambda(r_1, r_2 : \text{Ref}).\{\}$, where a function is represented in the lambda calculus style. Some constraints have parameters other than mere references, which are treated as variables inside, e.g. parallel distance in List 1. Third, we define the sketch concepts as \mathbb{L}^1 types composed of \mathbb{L}^0 types. To be specific, a composite type $\mathbf{T}_i^1 \in \mathbb{L}^1$ is a function with arguments $[\alpha_k]$ and members $t_{i,j}^0 : \mathbf{T}_j^0 \in \mathbb{L}^0$, which are connected through a composition operator $R_{\mathbf{T}_i^1} = \{p(q) | p, q \in [t_{i,j}^0] \cup [\alpha_k]\}$ that specifies how each pair of primitive elements binds together. For example, a coincident constraint $p = \lambda(r_1, r_2).\{\}$ may take a line primitive q as its first argument and bind to an argument α_k of the composite type as its second argument, i.e. $p(q, \alpha_k) \in R_{\mathbf{T}_i^1}$;

List 1: A domain-specific language formulating CAD sketch concepts

```

// Basic data types
Length, Angle, Coord, Ref
//  $\mathbb{L}^0$  primitive types
Line  $\rightarrow c_{start\_x}, c_{start\_y}, c_{end\_x}, c_{end\_y} : \text{Coord}$ 
Circle  $\rightarrow c_{center\_x}, c_{center\_y} : \text{Coord}, l_{radius} : \text{Length}$ 
...
//  $\mathbb{L}^0$  constraint types
Coincident  $\rightarrow \lambda(r_1, r_2 : \text{Ref}).\{\}$ 
Parallel Distance  $\rightarrow \lambda(r_1, r_2 : \text{Ref}).\{l_{dist} : \text{Length}\}$ 
...
//  $\mathbb{L}^1$  composite types
 $\mathbf{T}_i^1 \rightarrow \lambda([\alpha_k : \text{Ref}]).\{t_{i,j}^0 : \mathbf{T}_j^0 \in \mathbb{L}^0, R_{\mathbf{T}_i^1}([t_{i,j}^0] \cup [\alpha_k])\}$ 
// Sketch decomposition
 $S \rightarrow \{t_i^1 : \mathbf{T}_i^1 \in \mathbb{L}^1, R_S([t_i^1])\}$ 

```

78 on the other hand, an argument α_k may bind to a primitive q , which is specified by $\alpha_k(q) \in R_{\mathbf{T}_i^1}$.
79 Finally, an input sketch S is restructured as a collection of composite types $t_i^1 : \mathbf{T}_i^1 \in \mathbb{L}^1$, as well as
80 their connections specified by a corresponding composition operator R_S . R_S records how different
81 concepts bind through their arguments, which further transfers to \mathbb{L}^0 typed elements inside the
82 concepts and translates into the raw constraint relationships of the sketch graph.

83 Given the explicit formulation of CAD sketches through a DSL, the discovery of sketch concepts
84 becomes the task of learning program libraries \mathbb{L}^1 by induction on many sketch samples. Therefore,
85 our task resembles shape program synthesis that aims at building modular programs for generating
86 shapes [5, 7], and differs from works that use autoregressive language models to generate CAD sketch
87 programs one token at a time [6, 13, 17]. In Sec. 6.2, we show that the structured learning of CAD
88 sketches enables more robust auto-completion than unstructured language modeling.

89 The search of structured concepts is clearly a combinatorial problem with exponential complexity,
90 which is intractable unless we can exploit the inherent patterns in large-scale sketch datasets. However,
91 to enable deep learning based detection and search of structured concepts, we need to bridge the
92 implicit deep representations and the explicit and interpretable structures, which we build through the
93 following two key observations:

- 94 • **A concept has dual representations:** implicit and explicit. The implicit representation as
95 embeddings in latent spaces is compatible with deep learning, while the explicit representa-
96 tion provides structures on which desired properties (e.g. modularity) can be imposed.
- 97 • **A concept is a parameterized structure.** A concept is a composite type with fixed modular
98 structure for interpretability, but the structure is always instantiated by assigning parameters
99 to its component primitives when the concept is found in a sketch.

100 3.1 Method overview

101 According to the two observations, we design an end-to-end sketch concept learning framework by
102 self-supervised induction on sketch graphs. As shown in Fig. 2, the framework has two main steps
103 before loss computation: a detection step that generates implicit representations of concepts making
104 up the input sketch, and an explicit generation step that expands the implicit concepts into concrete
105 structures on which self-supervision targets like reconstruction and modularity are applied.

106 Building on a state-of-the-art detection architecture [2], the detection module D takes a sketch S as
107 input and detects the modular concepts within it, i.e. $\{\mathbf{q}_i\} = D(S, \{\bar{\mathbf{q}}_i\})$, where the concepts are
108 represented implicitly as latent codes $\{\mathbf{q}_i\}$, and $\{\bar{\mathbf{q}}_i\}$ are a learnable set of concept instance queries.
109 Notably, we apply vector quantization to the latent codes and obtain $\{\mathbf{q}'_i = \min_{\mathbf{p} \in \mathbb{L}^1} \|\mathbf{p} - \mathbf{q}_i\|_2\}$,
110 which ensures that each concept is selected from the common collection of learnable concepts \mathbb{L}^1
111 used for restructuring all sketches.

112 The explicit generation module is separated into two sub-steps, structure generation and parameter
113 instantiation, which ensures that the modular concept structures are explicit and reused throughout

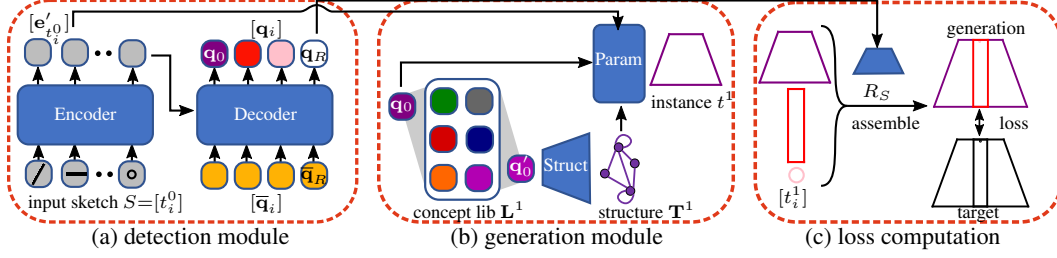


Figure 2: **Framework illustration.** (a) The detection module is a transformer network that detects from the sketch sequence $[t_i^0]$ implicitly encoded concepts $[q_i]$ and their composition q_R . (b) Each q is quantized against the concept library L^1 to obtain prototype q' , which is expanded by structure network into an explicit structure T^1 and further instantiated by parameter network into t^1 . (c) The collection of $[t_i^1]$ are assembled by composition operator R_S generated from q_R to obtain the final generated sketch graph, which is compared with the input sketch for loss computation.

different sketch instances. Specifically, the structure network takes each quantized concept code q'_i and generates its explicit form T_i^1 in terms of primitives and constraints of \mathbb{L}^0 types along with the composition operator $R_{T_i^1}$. Subsequently, the parameter network instantiates the concept structure by assigning parameter values to each component of T_i^1 conditioned on q_i and input sketch, to obtain t_i^1 . The composition operator R_S for combining $\{t_i^1\}$ is generated from a special latent code q_R transformed by D from a learnable token \bar{q}_R appended to $\{\bar{q}_i\}$.

The entire model is trained end-to-end by reconstruction and modularity objectives. In particular, we design loss functions that measure differences between the generated and groundtruth sketch graphs, in terms of both per-element attributes and pairwise references. Given our explicit modeling of encapsulated structures of the learned concepts, we can further enhance the modularity of the generation by introducing a bias loss that encourages in-concept references.

4 End-to-end sketch concept induction

4.1 Implicit concept detection

Sketch encoding A raw sketch S can be serialized into a sequence of \mathbb{L}^0 primitives and constraints. Previous works have adopted slightly different schemes to encode the sequence [6, 13, 17, 23, 24]. In this paper, we build on the previous works and take a simple strategy akin to [13, 24] for input sketch encoding. Specifically, we split each \mathbb{L}^0 typed instance t^0 into several tokens: *type*, *parameter*, and a list of *references*. For each of the token category, we use a specific embedding module. For example, parameters as scalars are quantized into finite bins before being embedded as vectors (see supplementary for the quantization details), and since there are at most five parameters for each primitive, we pack all parameter embeddings into a single code. On the other hand, each constraint reference as a primitive index is directly embedded as a code. Therefore, each token of a \mathbb{L}^0 typed instance is encoded as

$$e_{t^0.x} = \text{enc}_{type}(t^0) + \text{enc}_{pos}(t^0.x) + [\text{enc}_{param}(t^0.x) | \text{enc}_{ref}(t^0.x)],$$


where $t^0.x$ iterates over the split tokens, the type embedding is shared for all tokens of the instance, the position embedding counts the token index in the whole split-tokenized sequence of S , and parameter or reference embeddings are applied where applicable.

Concept detection We build the detection network as an encoder-decoder transformer following [2]. The transformer encoder operates on the sketch encoded sequence $[e_{t_i^0 \in S}]$ and produces the contextualized sequence $[e'_{t_i^0 \in S}]$ through layers of self-attention and feed-forward. The transformer decoder takes a learnable set of concept queries $[\bar{q}_i]$ of size k_{qry} plus a special query \bar{q}_R for composition generation, and applies interleaved self-attention, cross-attention to $[e'_{t_i^0}]$ and feed-forward layers to obtain the implicit concept codes $[q_i]$ and q_R . The concept codes are further quantized into $[q'_i]$ by selecting concept prototypes from a library L^1 implicitly encoding \mathbb{L}^1 , before being expanded into explicit forms.

148 4.2 Explicit concept structure generation

149 **Concept structure expansion** Given a library code $\mathbf{q}' \in \mathbb{L}^1$ representing a type $\mathbf{T}^1 \in \mathbb{L}^1$, through
 150 an MLP we expand its explicit structure as a collection of codes $[t_i^0]$ representing the \mathbb{L}^0 type instances
 151 $[t_i^0]$ and a matrix representing the composition $R_{\mathbf{T}^1}$ of $[t_i^0]$ and arguments (cf. List 1).

152 We fix the maximum number of \mathbb{L}^0 type instances to k_{L^0} (12 by default), and
 153 split the arguments into two groups, *inward arguments* and *outward arguments*,
 154 each of maximum number k_{arg} (2 by default). Each type code t_i^0 is decoded
 155 into discrete probabilities over \mathbb{L}^0 with an additional probability for null type
 156 ϕ to indicate the emptiness of this element (cf. Sec. 5.1), by $\text{dec}_{type}(\cdot)$ as
 157 the inverse of $\text{enc}_{type}(\cdot)$ in Sec. 4.1. An inward argument only points to a
 158 primitive inside the concept structure and originates from a constraint outside,
 159 and conversely an outward argument only points to primitives outside and
 160 originates from a constraint inside the concept (see inset for illustration); the split into two groups
 161 eases composition computation, as discussed below.



162 The composition operator $R_{\mathbf{T}^1}$ is implemented as an assignment matrix $\mathbf{R}_{\mathbf{T}^1}$ of shape $(2k_{L^0} + k_{arg}) \times$
 163 $(k_{L^0} + k_{arg})$, where each row corresponds to a constraint reference or inward argument, and each
 164 column to a primitive or outward argument. The two-fold coefficient of constraint references comes
 165 from that any constraint we considered in the dataset [16] has at most two arguments. Each row is a
 166 discrete probability distribution such that $\sum_j \mathbf{R}_{\mathbf{T}^1}[i, j] = 1$, with the maximum entry signifying that
 167 the i -th constraint/outward argument refers to the j -th primitive/inward argument. We compute $\mathbf{R}_{\mathbf{T}^1}$
 168 by first mapping the concept code \mathbf{q}' to a matrix of logits in the shape of $\mathbf{R}_{\mathbf{T}^1}$, and then applying
 169 softmax transform for each row. Notably, we avoid the meaningless loops of an element referring
 170 back to itself, and inward arguments referring to outward arguments, by masking the diagonal blocks
 171 $\mathbf{R}_{\mathbf{T}^1}[2i:2i+2, i], i \in [k_{L^0}]$ and the argument block $\mathbf{R}_{\mathbf{T}^1}[2k_{L^0}:, k_{L^0}:]$ by setting their logits to $-\infty$.

172 **Cross-concept composition** Aside from references inside a concept, references across concepts are
 173 generated to complete the whole sketch graph. We achieve cross-concept references by argument
 174 passing (see inset above for illustration). In particular, we implement the cross-concept composition
 175 operator R_S as an assignment matrix \mathbf{R}_S of shape $(k_{qry} \cdot k_{arg}) \times (k_{qry} \cdot k_{arg})$ directly mapped from \mathbf{q}_R
 176 through an MLP. Similar to the in-concept composition matrix, each row of the cross-concept matrix
 177 is a discrete distribution such that $\sum_j \mathbf{R}_S[i, j] = 1$, with the maximum entry signifying that the
 178 $(i \bmod k_{arg})$ -th outward argument of the $\lfloor i/k_{arg} \rfloor$ -th concept instance refers to the $(j \bmod k_{arg})$ -th
 179 inward argument of the $\lfloor j/k_{arg} \rfloor$ -th concept instance.

180 The complete cross-concept reference is therefore the product of three transport matrices:

$$\mathbf{R}_{ref}[t_i^1, t_j^1] = \mathbf{R}_{t_i^1}[2k_{L^0}:, k_{L^0}:] \times \mathbf{R}_S[i \cdot k_{arg}:(i+1) \cdot k_{arg}, j \cdot k_{arg}:(j+1) \cdot k_{arg}] \times \mathbf{R}_{t_j^1}[2k_{L^0}:, :k_{L^0}],$$

181 where $\mathbf{R}_{ref}[t_i^1, t_j^1]$ is a block assignment matrix of shape $2k_{L^0} \times k_{L^0}$. Intuitively, $\mathbf{R}_{ref}[t_i^1, t_j^1]$
 182 specifies how constraints inside t_i^1 refers to primitives of t_j^1 throughout all possible paths crossing the
 183 arguments of two concepts.

184 Collectively, we denote the complete reference matrix for all pairs of generated \mathbb{L}^0 elements as \mathbf{R} of
 185 shape $(2k_{qry} \cdot k_{L^0}) \times (k_{qry} \cdot k_{L^0})$, which includes in-concept and cross-concept references.

186 4.3 Concept instantiation by parameter generation

187 Instantiating a concept structure requires assigning parameters to the components where applicable.
 188 Therefore, as shown in Fig. 2, the parameter generation network takes a concept structure \mathbf{T}^1 and
 189 its implicit instance encoding \mathbf{q} as input, and produces the specific parameters for each \mathbb{L}^0 typed
 190 instances inside the concept. In addition, as the parameters of generated instances are directly related
 191 to the input parameters of the raw sketch S , we find it improves convergence and accuracy by allowing
 192 the parameter network to attend to the input tokens.

193 We implement the parameter network as a transformer decoder in a similar way as [2]. The instance
 194 code \mathbf{q} is first expanded to k_{L^0} tokens by a small MLP, which are summed with $[t_i^0]$ token-wise to
 195 obtain the query codes. The parameter network then transforms the query codes through interleaved
 196 layers of self-attention, cross-attention to the contextualized input sequence $[e'_{t^0}]$, and feed-forward,
 197 before finally mapped to explicit parameters in the form of probabilities over quantized bins, through
 198 a decoding layer $\text{dec}_{param}(\cdot)$ that is inverse of $\text{enc}_{param}(\cdot)$ in Sec. 4.1.

5 Induction objectives

Without any given labels of concepts, we use the following objectives to supervise the inductive network training: sketch reconstruction, concept quantization, and modularity enhancement.

5.1 Reconstruction loss

As discussed in Sec. 4, an input sketch is restructured into a set of sketch concepts which are expanded into a graph of primitives and constraints; the generated sketch graph \tilde{S} is compared with the input sketch S for reconstruction supervision.

The comparison of generated and target graphs requires a one-to-one correspondence between elements of the two graphs, on which the graph differences can be measured. However, it is nontrivial to find such a matching, because not only are there variable numbers of elements in the two graphs, but also both *elements* and *references between elements* must be taken into account for matching. To this end, we build a cost matrix that measures the difference for each pair of generated and target elements, in terms of their attributes and references, and apply linear assignment matching on the cost matrix [11, 8] to establish the optimal correspondence between two graphs.

Cost matrix construction To compare each pair of generated element and target element, we measure their type differences, and further use type casting to interpret the generated element as the target type, so that their parameters can be compared. To account for reference differences between the two elements, we compare the reference arrows by the differences of their pointed primitives.

Specifically, given the target graph S of k_{tgt} elements and the generated graph \tilde{S} of $k_{qry} \cdot k_{L^0}$ elements, we build the cost matrix \mathbf{C} of shape $k_{tgt} \times (k_{qry} \cdot k_{L^0})$ in two steps. First, for a pair of target element p and generated element q , we compare their type and parameter differences by cross-entropy. We denote the cost matrix in this stage as \mathbf{C}_{ury} , as it accounts for the element-wise unary distances between two graphs. Second, to measure the binary distances of references, for each target constraint element p and its r -th referenced primitive $p.r$, its distance from the generated references of element q is computed as (also illustrated by inset figure):

$$\mathbf{C}_{bry}[p, q] = \sum_{r \in \{0,1\}} \sum_{j \in \tilde{S}} \mathbf{R}[2q + r, j] \times \mathbf{C}_{ury}[p.r, j], \quad (1)$$


where $\mathbf{R}[2q + r, j]$ is the probability of q taking j as its r -th reference, as predicted by the composition operation (Sec. 4.2). Intuitively, the binary cost is a summation of unary costs weighted by predicted reference probabilities, where the unary costs measure how different a generated pointed primitive is from the target pointed primitive. The complete cost matrix is $\mathbf{C} = w_{ury} \mathbf{C}_{ury} + w_{bry} \mathbf{C}_{bry}$, with $w_{ury} = 50$, $w_{bry} = 1$; we give a larger weight to the unary costs because meaningful binary costs depend on reliable unary costs in the first place, as evident in Eq. (1).

Matching and reconstruction loss Given the cost matrix \mathbf{C} , we apply linear assignment to obtain a matching $\sigma : \tilde{S} \rightarrow S \cup \{\phi\}$ between \tilde{S} and S . Note that the number of elements of these two graphs can be different, but we have chosen k_{qry}, k_{L^0} such that the generated elements always cover the target elements. Therefore, $\sigma(q) = p$ assigns a matched generation $q \in M \subset \tilde{S}$ to a target $p \in S$, but assigns the rest unmatched generations $M' = \tilde{S} \setminus M$ to the empty target ϕ , i.e. $\sigma(M') = \phi$. The loss terms for matched generations are simply the corresponding cost terms $\mathbf{C}[\sigma(q), q]$, $q \in M$; for unmatched generations, we supervise its type to be the empty type ϕ and neglect its parameters or references. We denote the average loss of all generated terms as \mathcal{L}_{recon} .

Besides matching cost, we also use an additional reference loss to encourage the generated references to be sharp (i.e., \mathbf{R} being closer to binary). This loss complements the binary costs mentioned above by making sure that even if the generated primitives are similar, a generated constraint only refers to one primitive sharply. We define the sharp reference loss as

$$\mathcal{L}_{sharp} = -\frac{1}{|S_c|} \sum_{p \in S_c, r} \log \mathbf{R}[2\sigma^{-1}(p) + r, \sigma^{-1}(p.r)],$$

where p iterates over the target constraints S_c , $\sigma^{-1}(\cdot)$ is the inverse mapping from target element to generation, and we skip a term if $p.r$ does not exist for constraints with one reference.



246 5.2 Concept quantization loss

247 Following [20, 15], we optimize the concept code quantization against library \mathbf{L}^1 with:

$$\mathcal{L}_{vq} = \frac{1}{k_{qry}} \sum_{i \in [k_{qry}]} \|\text{sg}(\mathbf{q}_i) - \mathbf{q}'_i\| + \beta \|\mathbf{q}_i - \text{sg}(\mathbf{q}'_i)\|,$$

248 where $\text{sg}(\cdot)$ is the stop gradient operation. For training stability, we follow [15] and replace the first
 249 term with EMA updates of $\mathbf{q}' \in \mathbf{L}^1$. Furthermore, we improve spare code usage by reviving unused
 250 code in \mathbf{L}^1 periodically [15] (please refer to supplementary for details).

251 5.3 Modularity enhancement loss

252 We look for modular \mathbb{L}^1 concepts that have rich and meaningful encapsulated structures, rather than
 253 arbitrary groups of \mathbb{L}^0 elements that rely on cross-group references to recover the graph structures.
 254 This modularity can be enhanced by limiting the use of arguments for sketch concepts. Instead
 255 of allocating very few arguments as a hard constraint, we introduce a soft bias loss to encourage
 256 the restrictive use of arguments, which may still cover cases when more arguments are needed for
 257 accurate reconstruction. To be specific, we penalize the accumulated probability of elements pointing
 258 to outward arguments:

$$\mathcal{L}_{bias} = \frac{1}{|S_c|} \sum_{p \in S_c, r} \sum_{i \in [k_{arg}]} \mathbf{R}_{\mathbf{T}^1 \ni \sigma^{-1}(p)} [2\sigma^{-1}(p) + r, i + k_{L^0}],$$

259 where again p iterates over the target constraints S_c , $\sigma^{-1}(\cdot)$ is the inverse mapping from target
 260 element to generation, and $i + k_{L^0}$ slices the reference probabilities to outward arguments.

261 5.4 Total loss

262 The training objective sums up losses of reconstruction, concept quantization and modularity bias:

$$\mathcal{L}_{total} = w_{recon} \mathcal{L}_{recon} + w_{sharp} \mathcal{L}_{sharp} + w_{vq} \mathcal{L}_{vq} + w_{bias} \mathcal{L}_{bias},$$

263 where we empirically use weights $w_{recon} = 1, w_{sharp} = 20, w_{vq} = 1, w_{bias} = 25$ throughout all
 264 experiments unless otherwise specified in the ablation studies.

265 6 Results

266 **Dataset and implementation** Following previous works [6, 13, 17], we adopt the SketchGraphs
 267 dataset [16] which contains millions of real-world CAD sketches for training and evaluation. We
 268 filter the data by removing trivially simple sketches and duplicates, and limit the sketch complexity
 269 such that the number of primitives and constraints is within [20, 50]. As a result, we obtain around 1
 270 million sketches and randomly split them into 950k for training and 50k for testing. We defer network
 271 details to the supplementary and will open-source code and data to facilitate future research¹.

272 **Evaluation metrics** We evaluate the generated sketches in terms of reconstruction accuracy and
 273 sketch concept modularity, which are the two major objectives of our task. We measure reconstruction
 274 accuracy by the F-scores of generated primitives and constraints, where F-score is simply the harmonic
 275 mean of precision and recall. A generated primitive is considered a correct match with ground-truth
 276 if its type and parameters are correct, where for the scalar parameters we allow a threshold of 10% of
 277 quantization levels. A constraint is correct if and only if its type, parameter and references match
 278 ground-truth, i.e., the generated q is correct w.r.t target p iff q has the same type and parameters with
 279 p and the primitives $q.r$ and $p.r$ are correctly matched. Modularity is measured by the percentage of
 280 constraints with references entirely within the encapsulating concepts, among all correct constraints.

¹url to code and data will be updated soon.

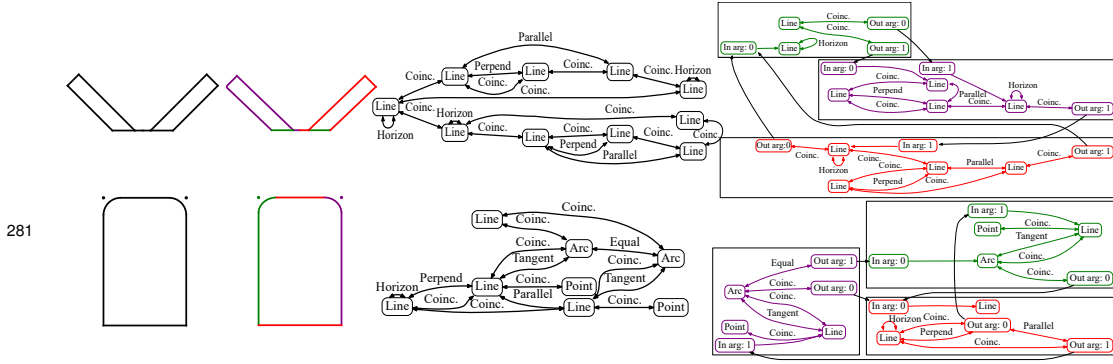


Figure 3: **Design intent parsing.** Left: input raw sketches and sketches restructured with concepts. Right: raw constraint graphs and modular constraint graphs. Primitives and constraints in the restructured sketches and graphs are colored according to concepts.




Figure 4: **Instances of concepts.** Left: two learned rectangle concepts with subtly different structures. Right: four sketches containing instances of these two concepts.




Figure 5: **Auto completion.** Each example shows the input partial sketch (black) and groundtruth completion (red), result of the autoregressive baseline, and our result (colored by concepts).




Figure 6: **Auto completion comparison.** Plotted are F-scores at different ratios of partial input.

Config	Primitive	Constraint	Modular(%)
$-C_{bry}$	0.993	0.290	34.2
$-\mathcal{L}_{sharp}$	0.983	0.104	100
$-\mathcal{L}_{bias}$	0.992	0.743	13.3
Ours	0.994	0.766	50.8

Table 1: **Loss ablation.** F-scores are reported for primitives and constraints.

6.1 Design intent interpretation

By training our model on the raw sketches with self-supervised induction losses, we obtain a result library of sketch concepts and a model for design intent parsing that interprets a given sketch into modular concepts and their combination. Indeed, we find the automatically discovered concepts capture natural design intents and modular structures. For example, through the restructured sketches and constraint graphs in Figs. 1 and 3, we find that our network decomposes sketches into modular structures like rectangles, line-arcs and parallel lines that align symmetrically, even though no such prior knowledge is applied during training except for concept modularity. Fig. 4 shows that a given concept can be used repetitively in different sketches, and structures with subtle differences in constraint relations can be detected and distinguished into different concepts of the library. Note that these subtle structural differences are subsumed in the input sketch graph, which makes them more difficult to detect. We refer to the supplementary for more examples of design intent parsing and instantiation of learned concepts, as well as quantitative analysis of the learned library.

6.2 Auto completion

Auto-completion is a critical feature of CAD modeling software for assisting designers. Given a partial sketch of primitives and their constraints, auto-completion aims at complementing them with the rest primitives and constraints to form regular and well-structured designs. Therefore, our concept detection and generation approach would naturally enhance the auto-completion task with better regularity. For training and evaluation, following previous work [17], we synthesize the partial input by removing a suffix of random length (up to 50%) from the sketch sequence, along with constraints that refer to the removed primitives, and make the model learn to generate the full sketches.

State-of-the-art methods [6, 13, 17, 23] formulate auto-completion through a combination of primitive and constraint generation models, both of which operate in an autoregressive fashion, with the constraint model conditioned on and referring (by pointers [22]) to the generated primitives. Since these works use diverse sketch encodings and have no publicly released code at submission time, for fair comparison, we implement the autoregressive baseline with our sketch encoding (Sec. 4.1).

Fig. 6 compares our method with autoregressive baseline under various primitive mask ratios: our method has superior primitive and constraint accuracy than the autoregressive baseline at almost all mask ratios. This difference confirms that since our method completes sketches concept-by-concept instead of primitive-by-primitive, more meaningful structures are likely to be generated. Our model also gains advantage by taking primitives and constraints together as input and generating primitives and constraints simultaneously, while in comparison the autoregressive baseline separates generation in two steps (primitives followed by constraints). Indeed, in practice CAD designers rarely finish all primitives first before supplementing the constraints, but rather apply constraints on partial primitives immediately whenever they form a design intent. Fig. 5 shows how our approach interprets the partial inputs and completes with modular concepts (see supplementary for more examples); in comparison, the autoregressive baseline does not provide such interpretable or regular completions.

6.3 Ablation study

To evaluate the impact of different loss terms of the induction objective (Sec. 5), we train several models in the absence of these losses respectively on the auto-encoding task. The results are shown in Table. 3. We see that removing the binary costs C_{bry} from reconstruction loss results in significant drop of constraint reconstruction, showing its necessity for constraint reference learning. Removing sharp reference loss \mathcal{L}_{sharp} similarly fails constraint reference learning, although modularity enhancement bias loss makes all constraint references inside concepts. Removing the modularity enhancement bias loss \mathcal{L}_{bias} only results in a slight drop in reconstruction quality but a significant drop in modularity, since without it cross-concept reference through arguments is more likely and therefore modularity suffers. We provide more ablation tests on hyper parameters like the numbers of concept queries k_{qry} and arguments k_{arg} in the supplementary.

7 Conclusion

CAD sketch concepts are meta-structures containing primitives and constraints that define modular sub-graphs and capture design intents. By formulating the sketch concepts as program libraries of a DSL, we present an end-to-end approach for discovering CAD sketch concepts through library induction learning. Key to our approach are the implicit-explicit representation of concepts and the separated structure generation and parameter instantiation for concept generation, which together enable the end-to-end training under self-supervised induction objectives. By training on large-scale sketch dataset, our approach enables the discovery of repetitive and modular concepts from raw sketches, and more structured and interpretable auto-completion than baseline autoregressive models.

Limitations and future work Design intents can be hierarchical [10], meaning that higher order meta-structures can be built out of lower order ones. In this sense, our framework only addresses the first order library induction, and should be extended for higher order library learning; toward this goal, we believe a progressive approach like [5] can be used with our framework as the one-step induction. In addition, similar strategies of end-to-end induction learning can be applied to constraint graphs involving 3D CAD operations or even more general programs in other domains, as long as they have similar declarative and parametric structures as sketch graphs.

References

- [1] Jorge D Camba, Manuel Contero, and Pedro Company. Parametric cad modeling: An analysis of strategies for design reusability. *Computer-Aided Design*, 74, 2016.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020.
- [3] Kevin Ellis, Daniel Ritchie, Armando Solar-Lezama, and Josh Tenenbaum. Learning to infer graphics programs from hand-drawn images. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [4] Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. Write, execute, assess: Program synthesis with a repl. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [5] Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sablé-Meyer, Lucas Morales, Luke Hewitt, Luc Cary, Armando Solar-Lezama, and Joshua B. Tenenbaum. Dreamcoder: Bootstrapping inductive program synthesis with wake-sleep library learning. In *Proceedings of the ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2021*, 2021.
- [6] Yaroslav Ganin, Sergey Bartunov, Yujia Li, Ethan Keller, and Stefano Saliceti. Computer-aided design as language. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [7] R. Kenny Jones, David Charatan, Paul Guerrero, Niloy J. Mitra, and Daniel Ritchie. Shapemod: Macro operation discovery for 3d shape programs. *ACM Transactions on Graphics (Siggraph)*, 40(4), 2021.
- [8] Harold W. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [9] Sofia Kyratzi and Philip Azariadis. A constraint-based framework to recognize design intent during sketching in parametric environments. *Computer-Aided Design and Applications*, 18(3), 2021.
- [10] Sofia Kyratzi and Philip Azariadis. Integrated design intent of 3d parametric models. *Computer-Aided Design*, 146, 2022.
- [11] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.
- [12] Jeffrey Otey, Pedro Company, Manuel Contero, and Jorge D. Camba. Revisiting the design intent concept in the context of mechanical cad education. *Computer-Aided Design and Applications*, 15(1), 2018.
- [13] Wamiq Para, Shariq Bhat, Paul Guerrero, Tom Kelly, Niloy Mitra, Leonidas J Guibas, and Peter Wonka. Sketchgen: Generating constrained cad sketches. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [14] PTC Inc. Onshape, 2022. URL <https://www.onshape.com/en/>. Accessed Apr. 20, 2022.
- [15] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [16] Ari Seff, Yaniv Ovadia, Wenda Zhou, and Ryan P. Adams. SketchGraphs: A large-scale dataset for modeling relational geometry in computer-aided design. In *ICML 2020 Workshop on Object-Oriented Learning*, 2020.
- [17] Ari Seff, Wenda Zhou, Nick Richardson, and Ryan P Adams. Vitruvion: A generative model of parametric CAD sketches. In *International Conference on Learning Representations*, 2022.
- [18] Yonglong Tian, Andrew Luo, Xingyuan Sun, Kevin Ellis, William T. Freeman, Joshua B. Tenenbaum, and Jiajun Wu. Learning to infer and execute 3d shape programs. In *International Conference on Learning Representations*, 2019.

- [19] Lazar Valkov, Dipak Chaudhari, Akash Srivastava, Charles Sutton, and Swarat Chaudhuri. Houdini: Lifelong learning as program synthesis. In *International Conference on Neural Information Processing Systems*, 2018.
- [20] Aaron van den Oord, Oriol Vinyals, and koray kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- [22] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- [23] Karl D.D. Willis, Pradeep Kumar Jayaraman, Joseph G. Lambourne, Hang Chu, and Yewen Pu. Engineering sketch generation for computer-aided design. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2021.
- [24] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *IEEE International Conference on Computer Vision (ICCV)*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [\[Yes\]](#)
 - (b) Did you describe the limitations of your work? [\[Yes\]](#) Please see Sec. 7.
 - (c) Did you discuss any potential negative societal impacts of your work? [\[Yes\]](#) Please see Appendix.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [\[Yes\]](#)
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [\[N/A\]](#) We present a deep learning algorithm rather than theorems in this paper.
 - (b) Did you include complete proofs of all theoretical results? [\[N/A\]](#)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [\[No\]](#) We will update the url in Sec. 6 to open-source our code and data soon.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) Please refer to Sec. 6 and Appendix.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[No\]](#) We did not report error bars due to lack of computational resources for running all comparisons and ablations multiple times consistently. But we note that the contrasts between different methods are very obvious, and we did run different methods in non-uniform repetitions and have found the same conclusions.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[Yes\]](#) Please see Appendix.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#) Please see Sec. 6.
 - (b) Did you mention the license of the assets? [\[Yes\]](#) The citation [16] clearly states that the SketchGraphs dataset we use has an MIT license.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[No\]](#) We will update the url in Sec. 6 to released source code and cleaned data soon.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [\[Yes\]](#) The MIT license of [16] permits research purpose usage as done in this paper.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[N/A\]](#) The CAD sketch graph data used in this paper is not personally related.

- 449 5. If you used crowdsourcing or conducted research with human subjects...
- 450 (a) Did you include the full text of instructions given to participants and screenshots, if
- 451 applicable? [N/A]
- 452 (b) Did you describe any potential participant risks, with links to Institutional Review
- 453 Board (IRB) approvals, if applicable? [N/A]
- 454 (c) Did you include the estimated hourly wage paid to participants and the total amount
- 455 spent on participant compensation? [N/A]