
Unsupervised Program Synthesis for Images By Sampling Without Replacement

Chenghui Zhou¹

Chun-Liang Li¹

Barnabás Póczos¹

¹Machine Learning Department, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA
{chenghuz, chunli1, bapoczos}@cs.cmu.edu

Abstract

Program synthesis has emerged as a successful approach to the image parsing task. Most prior works rely on a two-step scheme involving supervised pretraining of a Seq2Seq model with synthetic programs followed by reinforcement learning (RL) for fine-tuning with real reference images. Fully unsupervised approaches promise to train the model directly on the target images without requiring curated pretraining datasets. However, they struggle with the inherent sparsity of meaningful programs in the search space. In this paper, we present the first unsupervised algorithm capable of parsing constructive solid geometry (CSG) images into context-free grammar (CFG) without pretraining via non-differentiable renderer. To tackle the *non-Markovian* sparse reward problem, we combine three key ingredients—(i) a grammar-encoded tree LSTM ensuring program validity (ii) entropy regularization and (iii) sampling without replacement from the CFG syntax tree. Empirically, our algorithm recovers meaningful programs in large search spaces (up to 3.8×10^{28}). Further, even though our approach is fully unsupervised, it generalizes better than supervised methods on the synthetic 2D CSG dataset. On the 2D computer aided design (CAD) dataset, our approach significantly outperforms the supervised pretrained model and is competitive to the refined model.

1 INTRODUCTION

Image generation is extensively studied in machine learning and computer vision literature. Vast numbers of papers have explored image generation through low-dimensional latent representations [Goodfellow et al., 2014, Arjovsky et al., 2017, Li et al., 2017, Kingma and Welling, 2013, van den

Oord et al., 2017, Oord et al., 2016]. However, it is challenging to learn disentangled representations that allow control over each component of the generative models separately [Higgins et al., 2017, Kim and Mnih, 2018, Locatello et al., 2018, Chen et al., 2016]. In this paper, we tackle the problem of CFG program generation from constructive solid geometry (CSG) [Hubbard, 1990] and computer aided design (CAD) images, which are commonplace in engineering and design applications. Parsing a geometric image into a CFG program not only enables selective manipulations of the desired components while preserving the rest, but also provides a human-readable alternative to the opaque low-dimensional representations generated by neural networks. Our model for extracting the programs can be viewed as an encoder and the renderer that reconstructs the image as a decoder. We focus on a non-differentiable renderer (most design software are non-differentiable, e.g. Blender and Autodesk). They are more common than differentiable ones [Li et al., 2018, Liu et al., 2019, Kania et al., 2020], but they are also more challenging to work with neural networks because their discrete nature cannot be integrated into the network and the gradients w.r.t. the rendered images are *inaccessible*.

A common scheme for parsing images (e.g. CSG images) into programs (e.g. CFG programs) for non-differentiable renderers involves two steps: first use synthetic images with ground-truth programs for supervised pretraining, followed by REINFORCE fine-tuning [Sharma et al., 2018, Ellis et al., 2019] on the target image dataset. Sampling programs from a grammar can provide data suitable for supervision if the target images are restricted to combinations of geometric primitives specified in the grammar. There are two limitations of supervised methods. Firstly, it maximizes the likelihood of a single reference program while penalizing many other correct programs [Bunel et al., 2018] using maximum likelihood estimation (MLE). This observation is known as *program aliasing* and it adversely affects supervised learning’s performance. Secondly, it does not generalize well to the test images not generated by the grammar. REINFORCE fine-tuning is proposed to remedy the two limitations [Bunel

et al., 2018, Sharma et al., 2018]. The transition between the supervised pretraining and REINFORCE fine-tuning is delicate, however, because the model is sensitive to bad gradient updates that cause the grammar structure to collapse among the generated programs. Additionally, the quality of the curated pretraining dataset can limit the downstream model’s ability to generalize.

In this paper, however, we focus on the more interesting and more challenging *unsupervised task*, when ground-truth programs of the images are not available for training. Compared with the supervised pretrained alternative, unsupervised image parsing is under explored. The benefits of an unsupervised approach are:

- **Train directly on the target domain.** Instead of training on a curated synthetic dataset that mimic the target dataset, e.g. combinations of geometric shapes specified in the CFG imitating a CAD dataset exceeding the primitives of the grammar, we can train directly on the target images. The synthetic dataset may not reflect the target image dataset accurately and can subsequently lead to insufficient generalization during test time.
- **Treat all correct programs equally.** When multiple programs correspond to the reconstruction of the same image, while supervised methods only optimize for programs included in the synthetic ground-truth dataset, the equivalent programs receive equal rewards under RL.

Despite the benefits, designing an approach without supervised pretraining is challenging. Because of the problem’s discrete nature, we rely on tools from reinforcement learning, such as REINFORCE. However, the program space grows exponentially with the length of the program and valid programs are too sparse in the search space to be sampled frequently enough to learn. Training with the naive REINFORCE provides no performance gain in our experiments. RL techniques such as Hindsight Experience Replay [Andrychowicz et al.] that mitigate the sparse reward problem cannot be applied here due to the Markov assumption on the model. We improve the sample efficiency of the REINFORCE algorithm and show that our improved approach achieves competitive results to a two-step model. We further demonstrate that our method generalizes better on a synthetic 2D CSG dataset than a supervised method. On a 2D CAD dataset, which was *NOT* generated by CFG and thus cannot be captured sufficiently by a synthetic dataset, our method exceeds the results of a pretrained model by a large margin and performs competitively to the refined models.

Here we summarize the key ingredients that help successfully learn to parse an image *without program supervision* while using a non-differentiable renderer *without direct gradient propagation* w.r.t. the rendered images:

- We incorporate a **grammar-encoded tree LSTM** to

impose a structure on the search space such that the algorithm is sampling a path in the CFG syntax tree top-down. This guarantees validity of the output program.

- We propose an **entropy estimator** suitable for sampling top-down from a syntax tree to encourage exploration of the search space by entropy regularization.
- Instead of relying on naive Monte Carlo sampling, we adopt **sampling without replacement** from the syntax tree to obtain better entropy estimates and REINFORCE objective for faster convergence.

2 RELATED WORK

Program synthesis attracts growing interests from researchers in machine learning. Supervised training is a natural choice for input/output program synthesis problems [Parisotto et al., 2016, Chen et al., 2018a, Devlin et al., 2017, Yin and Neubig, 2017, Balog et al., 2017, Zohar and Wolf, 2018]. Shin et al. [2018] use the input/output pairs to learn the execution traces. Bunel et al. [2018] use RL to address program aliasing, however, supervised pretraining is still necessary to reap its benefits. Approaches to ensure valid outputs involve syntax checkers [Bunel et al., 2018] or constructions of abstract syntactic trees (AST) [Parisotto et al., 2016, Yin and Neubig, 2017, Kusner et al., 2017, Chen et al., 2018b]. A graph can also model the information flow in a program [Brockschmidt et al., 2018].

Vision-as-inverse-graphics focuses on parsing a scene into a collection of shapes or 3D primitives, e.g. cars or trees, with parameters, e.g. colors or locations, that imitates the original scene [Tulsiani et al., 2017, Romaszko et al., 2017, Wu et al., 2017]. Yao et al. [2018] further manipulate the objects de-rendered, such as color changes. Stroke-based rendering creates an image similar to how we write and draw. Some of the examples are recreating paintings by imitating a painter’s brush strokes [Huang et al., 2019], drawing sketches of objects [Ha and Eck, 2017]. SPIRAL [Ganin et al., 2018] is an adversarially trained deep RL agent that can recreate MNIST digits and Omniglot characters. Contrary to our problem, a grammar structure is unnecessary to both vision-as-inverse-graphics and stroke-based rendering.

Research on converting images to programs relates more closely to our work [Sharma et al., 2018, Ellis et al., 2019, 2018, Liu et al., 2018, Shin et al., 2019, Beltramelli, 2018, Kania et al., 2020]. Tian et al. [2019], Kania et al. [2020] incorporate differentiable renderers into the learning pipeline while we treat our renderer as an external process independent from the learning process, thus unable to propagate gradient through the renderer. Furthermore, Kania et al. [2020] construct the parse tree bottom up and pre-determines the number of leaves, as opposed to top down like ours which is more general. Ellis et al. [2018] use neural networks to extract shapes from hand-drawn sketches, formulate the grammatical rules as constraints and obtain the final programs

by optimizing a constraint satisfaction problem. Similarly, Du et al. [2018] cast the problem of parsing a 3D model into a CSG tree as a constraint satisfaction problem and solve by an existing SAT solver. This process can be computationally expensive compared to neural network based solutions. More relevantly, Sharma et al. [2018] perform program synthesis by supervised pretraining before RL fine-tuning to generalize to a CAD dataset. Ellis et al. [2019] pretrain a policy with supervision from synthetic data and learns a value function with REINFORCE. Both are used for pruning unpromising candidates during test time. Their reward function is binary and cannot approximate images not generated by a grammar, as opposed to our model.

3 PROPOSED ALGORITHM

CSG Image and CFG Program. We use constructive solid geometry (CSG) [Hubbard, 1990] to describe an image. The input of our model are images constructed from geometric shapes (e.g. square, circle, ...) each with a designated size and location (see Figure 1). The outputs of the model are context-free grammar (CFG) programs. In the CFG specification [Sharma et al., 2018], S , T , and P are non-terminals for the start, operations, and shapes. The rest are terminals, e.g. $+$ (union), $*$ (intersection), $-$ (subtraction), and $c(48, 16, 8)$ stands for a circle with radius 8 at location (48, 16) in the image. Figure 3 contains examples of CSG images and their corresponding programs. Each line below is a *production rule* or just *rule* for simplicity:

$$S \rightarrow E \quad (1)$$

$$E \rightarrow EET|P \quad (2)$$

$$T \rightarrow + | - | * \quad (3)$$

$$P \rightarrow SHAPE_1 | SHAPE_2 | \dots | SHAPE_n. \quad (4)$$

3.1 LEARNING WITH REINFORCE

A model trained with the REINFORCE objective only [Sharma et al., 2018] is unable to improve beyond the lowest reward (Section 4.1.1). We propose three components to enable the RL approach to learn in this setting: (i) a tree model entropy estimator to encourage exploration; (ii) sampling without replacement in the program space to facilitate optimization and further encourage exploration; (iii) a grammar-encoded tree LSTM to ensure valid output sequences with an image stack to provide intermediate feedback. We start this section by discussing objective and reward function.

Objective Function Our model consists of a CNN encoder for input images, an embedding layer for the actions, and an RNN for generating the program sequences (see Figure 1). The model is trained with entropy regularized REINFORCE [Williams, 1992]. Here, let $\mathcal{H}(s)$ and $f(s)$ stand for the

entropy (we will define this later) and reward function of sequence s , respectively, and let θ denote the parameters of the model. The objective is optimized as follows

$$\Delta\theta \propto \mathbb{E}_{s \sim P_\theta(s)} [\nabla_\theta \log P(s) f(s)] + \alpha \nabla_\theta \mathcal{H}(s) \quad (5)$$

Reward Function The output program s is converted to an image \mathbf{y} by a non-differentiable renderer. The image is compared to the target image \mathbf{x} and receives a reward $f(s) = R(\mathbf{x}, \mathbf{y})$. We use Chamfer Distance (CD) as part of the reward function. The CD calculates the average matching distance to the nearest feature and is a greedy estimation of image similarity, unlike Optimal Transport (OT). However, OT is not computationally feasible for RL purpose.

Algorithm 1 Sampling w/o Replacement Tree LSTM

Input: Target Image \mathbf{x} , Number of samples k

Initialize: Grammar stack S , Image stack I , Sample set \mathbb{B}

Encode the target image $\tilde{T} \leftarrow \text{Encode}(T)$

$\mathbb{B} = \{s_0^i, G_{s_0^i}, \phi_{i,0} | s_0^i = \emptyset, G_{s_0^i} = 0, \phi_{i,0} = 0\}$
for $i \in 1, 2, \dots, k+1$ and $\mathcal{H}(v) = 0$

for $j := 1$ to n **do**

$\Phi_{:,j}, \mathcal{H}_{i,j} \leftarrow \text{TreeLSTM}(S, I, \mathbb{B}, \mathbf{x})$ (Algorithm 3)

$\mathbb{B} \leftarrow \text{Sample_w/o_Replacement}(\Phi_{:,j}, \mathbb{B}, k)$

(See Algorithm 2 and [Kool et al., 2019])

$\hat{\mathcal{H}}_D \leftarrow \hat{\mathcal{H}}_D + \sum_{j=1}^n \frac{1}{W_j(S)} \sum_{s^i \in S} \frac{p_\theta(s_j^i)}{q_{\theta,\kappa}(s_j^i)} \mathcal{H}_{i,j}$

(Equation 4 in Appendix B)

if $s_j^i \in \mathcal{G}$ **then** $S_i.\text{push}(s_j^i)$

else $I_i.\text{push}(s_j^i), \forall s^i \in \mathbb{B}$

end for

$\mathbf{y}_i = \text{Render}(s^i)$ for $i \in 1, 2, \dots, k$

Maximize $\mathbb{E}[\sum_{i=1}^k R(\mathbf{x}, \mathbf{y}_i)] + \alpha \hat{\mathcal{H}}_D$

Formally, let $x \in \mathbf{x}$ and $y \in \mathbf{y}$ be pixels in each image respectively. Then the distance $Ch(\mathbf{x}, \mathbf{y})$ is

$$Ch(\mathbf{x}, \mathbf{y}) = \frac{1}{|\mathbf{x}|} \sum_{x \in \mathbf{x}} \min_{y \in \mathbf{y}} \|x - y\|_2 + \frac{1}{\|\mathbf{y}\|} \sum_{y \in \mathbf{y}} \min_{x \in \mathbf{x}} \|x - y\|_2 \quad (6)$$

The CD is scaled by the length of the image diagonal (ρ) [Sharma et al., 2018] such that the final value is between 0 and 1. For this problem, the reward $1 - Ch(\mathbf{x}, \mathbf{y})$ mostly falls between 0.9 and 1. We exponentiate $1 - Ch(\mathbf{x}, \mathbf{y})$ to the power of $\gamma = 20$ to achieve smoother gradients [Laud, 2004]. We add another pixel intersection based component to differentiate shapes with similar sizes and locations. The final reward function is defined as:

$$R(\mathbf{x}, \mathbf{y}) = \max(\delta, (1 - \frac{Ch(\mathbf{x}, \mathbf{y})}{\rho})^\gamma + \frac{\sum_{x \in \mathbf{x} \cap \mathbf{y}} 1}{\sum_{x \in \mathbf{x}} 1}) \quad (7)$$

The first and second part of the reward function provide feedback on the physical distance and similarity of the prediction respectively. We clip the reward below $\delta = 0.3$ to

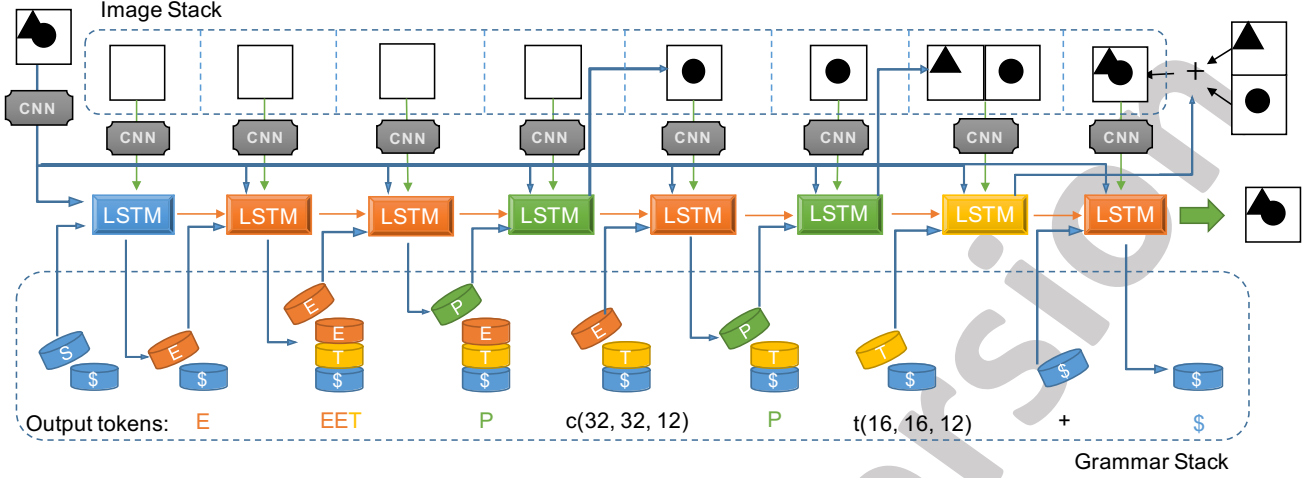


Figure 1: This is an example of the grammar encoded tree LSTM at work. The top layer of images demonstrates the image stack and the bottom layer demonstrates the grammar stack. The blue, orange, yellow and green colored LSTM cell generates grammatical tokens according to the CFG rule 1, 2, 3 and 4 respectively. In implementation, we can constrain the output space by adding a mask to the output of the LSTM and render the invalid options with close to zero probability of being sampled.

simplify it when the quality of the generated images are poor. A low reward value provides little insight on its performance and is largely dependent on its target image. Similar reward clipping idea was used in DQN [Mnih et al., 2013].

3.2 EXPLORATION WITH ENTROPY REGULARIZATION

Entropy regularization in RL is a standard practice for encouraging exploration. Here we propose an entropy estimation for sampling top-down from a syntax tree. Let S denote a random variable of possible programs. Its entropy is defined by $\mathcal{H}(S) = \mathbb{E}[-\log P(S)]$ ¹. Estimating $\mathcal{H}(S)$ is challenging because the possible outcomes of S is exponentially large and we cannot enumerate all of them. Given the distribution P , the naive entropy estimator is

$$\hat{\mathcal{H}} = -\frac{1}{K} \sum_{i=1}^K \log P(s^i), \quad (8)$$

where $\{s^i\}_{i=1}^K$ are iid samples from P . In practice, when K is not exponentially large, this estimator has huge variance. We provide an improved estimator designed for our syntax tree: First, we decompose the program S into $S = X_1 \dots X_n$, where each X_j is the random variable for the token at position j in the program. Under autoregressive models (e.g. RNN), we can access the conditional probabilities, and this allows us to construct decomposed entropy estimator $\hat{\mathcal{H}}_D$ as

$$\hat{\mathcal{H}}_D = \frac{1}{K} \sum_{i=1}^K \sum_{j=1}^n \mathcal{H}(X_j | X_1 = x_1^i, \dots, X_{j-1} = x_{j-1}^i), \quad (9)$$

where $s^i = x_1^i, \dots, x_n^i$, and $\mathcal{H}(X_j | X_1 = x_1^i, \dots, X_{j-1} = x_{j-1}^i)$ is the conditional entropy. The below lemma states that $\hat{\mathcal{H}}_D$ is indeed an improved estimator over $\hat{\mathcal{H}}$.

Lemma 3.1. *The proposed decomposed entropy estimator $\hat{\mathcal{H}}_D$ is unbiased with lower variance, that is $\mathbb{E}[\hat{\mathcal{H}}_D] = \mathcal{H}(S)$ and $\text{Var}(\hat{\mathcal{H}}_D) \leq \text{Var}(\hat{\mathcal{H}})$.*

The proof can be found in the Appendix C and D.

Algorithm 2 Sampling_w/o_Replacement

Input: Log probability at time j $\Phi_{:,j}$ Beam Set \mathbb{B} ,
Number of beams k

function SAMPLING_W/O_REPLACEMENT($\Phi_{:,j}, \mathbb{B}, k$)

$\tilde{\mathcal{G}}_j \leftarrow \emptyset$

for $s_{j-1}^i, G_{s_{j-1}^i} \in \mathbb{B}$ and $\vec{\phi}_{i,j} \in \Phi_{:,j}$ **do**

$\tilde{\mathcal{G}}_{\phi_{i,j}} \sim \text{Gumbel}(\vec{\phi}_{i,j})$

$Z_{i,j} \leftarrow \max(\tilde{\mathcal{G}}_{\phi_{i,j}})$

Calculate $\tilde{\mathcal{G}}_{\phi_{i,j}}$ (Equation 2 in Appendix B)

Aggregate the values in the vector $\tilde{\mathcal{G}}_{\phi_{i,j}}$

$\tilde{\mathcal{G}}_j \leftarrow \tilde{\mathcal{G}}_j \cup \tilde{\mathcal{G}}_{\phi_{i,j}}$

end for

Choose top $k+1$ values in $\tilde{\mathcal{G}}_j \in \mathbb{R}^{(k+1) \cdot A}$

and form the new beam set

$\tilde{\mathbb{B}} = \{s_j^i, G_{s_j^i}, \phi_j^i | s_j^i = \tilde{s}_{j-1}^i \cup s_j^i, G_{s_j^i} = \tilde{G}_{i,j}, \phi_{i,j} = \log P(s_j^i)\}$ where $i = 1, 2, \dots, k+1$

return $\tilde{\mathbb{B}}$

end function

¹Following Rule (1) and Rule (2) we overload S and P .




Figure 2: The left most image demonstrates the entropy value increases over 700 iterations by sampling 20 distinct samples with and without replacement as well as sampling 40 samples with replacement. The second image shows the initial distribution. The third and fourth images show the final distributions.

3.3 EFFECTIVE OPTIMIZATION BY SAMPLING WITHOUT REPLACEMENT

After establishing our REINFORCE method with entropy regularization objective, now we show the intuition behind choosing sampling without replacement (SWOR) over sampling with replacement (SWR). For this explanation, we use a synthetic example (Figure 2).

We initialize a distribution of $m = 100$ variables with three of them having significantly higher probability than the others (Figure 2 (2)). The loss function is entropy. Its estimator is $\frac{1}{m} \sum_{i=1}^m \log p_i$ for SWR and $\sum_{i=1}^m \frac{p_i}{q_i} \log p_i$ for SWOR. In both cases, p_i is the i -th variable’s probability. q_i is the re-normalized probability after SWOR. $\frac{p_i}{q_i}$ is the importance weighting. The increase in entropy by sampling 20 variables without replacement is more rapid than 40 variables with replacement. At the end of the 700 iterations, the distribution under SWOR is visibly more uniform than the other. SWOR would achieve better exploration than SWR.

To apply SWOR to our objective, the REINFORCE objective and the entropy estimator require importance weightings. Let s_j^i denotes the first j elements of sequence s^i :

$$\nabla_{\theta} \mathbb{E}_{s \sim p_{\theta}(s)} [f(s)] \approx \sum_{s^i \in S} \frac{\nabla_{\theta} p_{\theta}(s^i)}{q_{\theta}(s^i)} f(s^i) \quad \text{and,} \quad (10)$$

$$\hat{\mathcal{H}}_D \approx \sum_{j=1}^n \sum_{s^i \in S} \frac{p_{\theta}(s_j^i)}{q_{\theta}(s_j^i)} \mathcal{H}(X_j | X_1 = x_1^i, \dots, X_{j-1} = x_{j-1}^i) \quad (11)$$

Implementing SWOR on a tree structure to obtain the appropriate set of programs S is challenging. It is not practical to instantiate all paths and perform SWOR bottom-up. Instead, we adopt a form of stochastic beam search by combining top-down SWOR with Gumbel trick that is equivalent to SWOR bottom-up [Kool et al., 2019]. The sampling process is described in Algorithm 2. For more detailed explanation, the implementation of the re-normalized probability $q_{\theta}(s^i)$ as well as some additional tricks of variance reduction for the objective function, please refer to Appendix B.

3.4 GRAMMAR ENCODED TREE LSTM

We introduce a *grammar-encoded tree LSTM* which encodes the production rules into the model, thus guaranteed to generate correct programs, and significantly reduce the search space during the training [Kusner et al., 2017, Alvarez-Melis and Jaakkola, 2016, Parisotto et al., 2016, Yin and Neubig, 2017]. There are 3 types of production rules in the grammatical program generation – shape selection (P), operation selection (T), and grammar selection (E). Grammar selection in this problem setting includes $E \rightarrow EET$, and $E \rightarrow P$ and they decide whether the program would expand. We denote the set of shape, operations and non-terminal outcomes (e.g. EET in Rule (2)) to be \mathcal{P} , \mathcal{T} and \mathcal{G} respectively. A naive parameterization is to let the candidate set of the LSTM output to be $\{S, \$\} \cup \mathcal{T} \cup \mathcal{P}$, where $\$$ is the end token, and treat it as a standard language model to generate the program [Sharma et al., 2018]. The model does not explicitly encode grammar structures, and expect the model to capture it implicitly during the learning process. The drawback is that the occurrence of valid programs is sparse during sampling and it can prolong the training process significantly.

Algorithm 3 TreeLSTM Model

Input: Grammar Stack S , Image Stack I ,
Target Image x , Sample Set \mathbb{B}

function TREELSTM(S, I, \mathbb{B}, x)

$\tilde{x} \leftarrow \text{Encode}(x)$

for $s_{j-1}^i, \phi_{i,j-1} \in \mathbb{B}$ **do**

$g_i \leftarrow S_i.\text{pop}()$

$\tilde{g}_i \leftarrow \text{Embed}(g_i)$

$\tilde{I}_i \leftarrow \text{Encode}(I_i)$

$H_{i,j} \leftarrow \text{LSTM}(\tilde{g}_i, \tilde{I}_i, \tilde{x}, H_{i,j-1})$

$\mathbf{p}_{i,j} \leftarrow \text{softmax}(f(H_{i,j}) + \text{Mask}(g_i))$

Estimate entropy at this level:

$\mathcal{H}_{i,j} = \tilde{\mathbf{p}}_{i,j} \cdot \log \tilde{\mathbf{p}}_{i,j}$

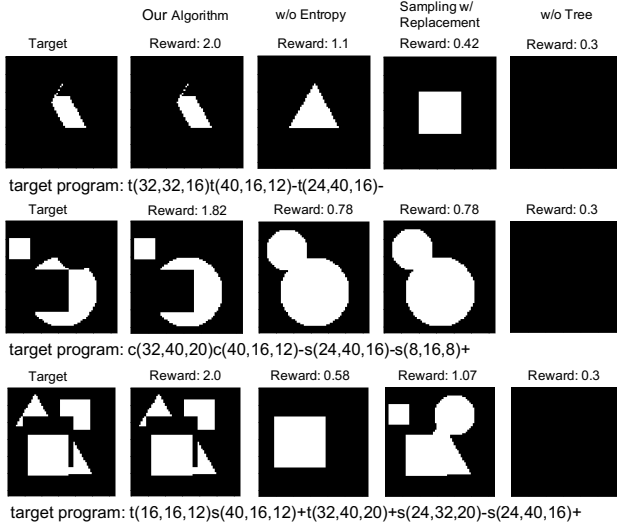
Update the log probabilities of partial sequences

$\tilde{\phi}_{i,j} = \tilde{\mathbf{I}} \cdot \phi_{i,j-1} + \log \mathbf{p}_{i,j}$

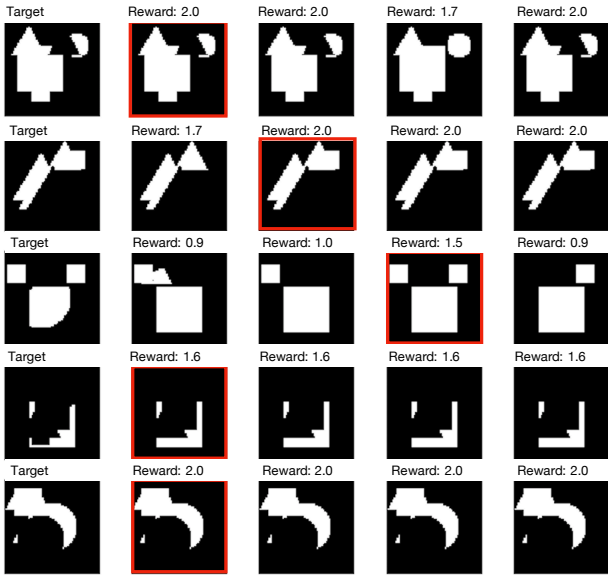
end for

return $\Phi_{:,j}, \mathcal{H}_{i,j}$

end function



(a) Example output of each method



(b) Example outputs of sampling without replacement

Figure 3: (a) We show a target image from each dataset and attach its correct program below. To the right are the reconstructed output programs from our algorithm and three variants each removing one design component. The reward is on top of the reconstructed images. (b) Some reconstructed example output programs of our algorithm. Each row represents one data point. The leftmost images of the five columns are the target images and the four columns to their right are the reconstructed outputs of four samples. The final output highlighted in red has the highest reward.

The proposed model can be described as an RNN model with a *masking mechanism* by maintaining a grammar stack to rule out invalid outputs. We increase the size of the total output space from $2 + |\mathcal{P}| + |\mathcal{T}|$ of the previous approach (e.g. [Sharma et al., 2018]) to $2 + |\mathcal{P}| + |\mathcal{T}| + |\mathcal{G}|$ by including the non-terminals. During the generation, we maintain a stack to trace the current production rule. Based on the current non-terminal and its corresponding expansion rules, we use the masking mechanism to weed out the invalid output candidates. Take the non-terminal T for example, we mask the invalid outputs to reduce the candidate size from $2 + |\mathcal{P}| + |\mathcal{T}| + |\mathcal{G}|$ to $|\mathcal{T}|$ only. In this process, the model will produce a sequence of tokens, including grammatical, shape and operation tokens. We only keep the terminals as the final output program and discard the rest. The resulting programs are ensured to be grammatically correct. During the generation process, grammatical tokens are pushed onto the grammar stack while intermediate images and operations are pushed onto an *image stack*. Images in the image stack are part of the input to the LSTM to aid the inference in the search space. A step-by-step guide through the tree LSTM for better understanding is in Appendix A and Figure 1 is a visual representation of the process.

4 EXPERIMENTS

There are two datasets we used for experiments – a synthetic dataset generated by the CFG specified in Section 3 and a 2D CAD furniture dataset. We compared the result of our algorithm to that of the same neural network trained with supervision of ground-truth programs. We observed that the supervised model shows poorer generalization in both datasets despite its access to additional ground-truth programs. We provide qualitative and quantitative ablation study of our algorithm on the synthetic dataset. We also showed that our algorithm can *approximate* the CAD images with programs despite not having exact matches. On the CAD furniture dataset, it outperforms the supervised pretrained model and achieves competitive result to the refined model. Additionally, we verify empirically that the stepwise entropy estimator (Equation 9) indeed has smaller variance than the naive estimator (Equation 8) as proven in Lemma 3.1.

4.1 SYNTHETIC DATASET STUDY

We use three synthetic datasets to test our algorithm. The action space includes 27 shapes (Figure 1), 3 operations and 2 grammar non-terminals to create a 64 by 64 images. The search space for an image up to 3 shapes (or program length 5) is around 1.8×10^5 and it gets up to 1.1×10^9 for 5 shapes (or program length 9). We separate our dataset by the length of the program to differentiate images with increasing complexity. Our synthetic dataset is generated by filtering out

Type	Length 5	Length 7	Length 9	2D CAD
Training set size	3600	4800	12000	10000
Testing set size	586	789	4630	3000

Table 1: Dataset statistics.

the duplicates and empty images in combinations of shape and operation actions in text. Images are considered duplicates if only 120 pixels are different between the two and are considered empty if there are no more than 120 pixels in the image. Table 1 contains the dataset size information.

4.1.1 Ablation Study of Design Components

For these 3 datasets, we sampled 19 programs without replacement for each target image during training. The negative entropy coefficient is 0.05 and the learning rate is 0.01. We use SGD with 0.9 momentum.

Removing either one of the three design components has reduced the performance of our algorithm. Under sampling with replacement setting, the model is quickly stuck at a local optimum (Figure 4 (yellow)). Without the entropy term in the objective function, the reward function is only able to improve on the length 5 dataset but fails to do so on longer programs. Both techniques have facilitated exploration that helps the model to escape local minimum. Without tree structure, the reward stays around the lowest reward (Figure 4 (green)) because the program is unable to generate valid programs. Grammar encoded tree LSTM effectively constrains the search space such that the sampled programs are valid and can provide meaningful feedback to the model.

We allow variations in the generated programs as long as the target images can be recovered, thus we evaluate the program quality in terms of the reconstructed image’s similarity to the target image. We measure our converged algorithm’s performance (Table 2) on the three test sets with Chamfer and IoU reward metrics (Equation 7 first and second term). The perfect match receives 1 in both metrics. Figure 3 provides some qualitative examples on the algorithms.

4.1.2 Comparison With Supervised Training

We compare a supervised learning method’s train and test results on the synthetic dataset, using the same neural network model as in the unsupervised method. The input at each step is the concatenation of embedded ground truth program and the encoded final and intermediate images. We use the same Chamfer reward metric as in Table 2 to measure the quality of the programs. The test results of the supervised method worsen with the increasing complexity (program length) while the train results are almost perfect across all three datasets. The unsupervised method receives consistently high scores and generalizes better to new data

Test Metric	Length 5	Length 7	Length 9
Chamfer Reward \uparrow	0.98	0.96	0.96
IoU Reward \uparrow	0.99	0.96	0.96

Table 2: The performance of the converged model of our algorithm on the test set measured with Chamfer reward and IoU reward.

in comparison to the supervised method (Table 3). This phenomenon can be explained by program aliasing [Bunel et al., 2018]. The RL method *treats all correct programs equally* and directly optimizes over the reward function in the image space while the supervised method is limited to the content of the synthetic dataset and only optimizes over the loss function in the program space.

4.1.3 Supervised Pretrained On Limited Data With REINFORCE Fine-tuning

In this experiment we pretrained the supervised model on a third of the synthetic training dataset till convergence. We take the model and further fine-tune it with vanilla REINFORCE on the full training sets. We report the reward throughout fine-tuning process (Figure 5) and it dropped sharply in all three datasets. Our explanation is that while the output of the original supervised pretrained model follow grammatical structures, they are not able to retain the structure consistently after updates during the refinement process, which leads to the collapse.

4.2 2D CAD FURNITURE DATASET STUDY

The dataset used in this experiment is a 2D CAD dataset [Sharma et al., 2018] that contains binary 64×64 images of various furniture items. We apply our algorithm to this problem with an action space of 396 basic shapes plus the operations and grammatical terminals described in Section 3. We limit the number of LSTM iterations to 24 steps, which corresponds to a maximum of 6 shapes. For an image up to 6 shapes the search space is 9.4×10^{17} . If we remove the grammar-encoded tree structure, the search space is 3.8×10^{28} . In order to scale up to such a big search space from the synthetic experiment, we increase the number of programs sampled without replacement to 550. The higher number usually corresponds to faster convergence and better performance at convergence but the performance gain diminishes at around 500 for this problem. The learning rate and entropy values used are 0.01 and 0.007 respectively. We train the model with only the Chamfer reward (first part of the Equation 7) because these images are not generated by CFG and exact matching solutions do not exist. During training, the reward converges to 0.72. Qualitative results are reported from the train and test set (Figure 6). The program reconstructions are able to capture the overall profile of the target images. However, the cutouts and angles deviate from the original because the shape actions consist




Figure 4: From left to right, we have reward per batch for programs of length 5, 7 and 9. It demonstrates the performance of our algorithm and controlled comparison in performance with alternative algorithms by removing one component at a time.




Figure 5: REINFORCE fine-tuning with pretrained model.

Chamfer Reward \uparrow	length 5	length 7	length 9
Training	1.00	1.00	0.99
Testing	0.99	0.91	0.83

Table 3: Supervised training results.

solely of unrotated squares, perfect circles and equilateral triangles.

4.2.1 Comparison With Supervised Pretraining

In this section, we measure the image similarity directly in Chamfer Distance (CD) (Equation 6) for comparison. We pretrained a model on 300k, 150k and 30k ground truth programs (including duplicates) each with learning rate of 0.001. We selected the pretrained models that reaches the lowest CD (at 1.41, 2.00, 2.79 respectively) on a synthetic validation set. We further fine-tuned the pretrained models on the CAD dataset with learning rate 0.006. The transition between pretraining and fine-tuning is delicate. The grammar structure of the output programs (Figure 5) collapses when we set the learning rate to be 0.01 (as opposed to 0.006) or fine-tune an unconverged pretrain model. Our

Chamfer Distance \downarrow	k = 1	k = 3	k = 5
30k Supervised Model	4.09	3.38	3.02
30k RL Refined Model	1.92	1.83	1.79
30k SWOR RL Refined Model	1.96	1.82	1.77
150k Supervised Model	3.64	2.89	2.63
150k RL Refined Model	1.91	1.79	1.73
150k SWOR RL Refined Model	1.93	1.79	1.73
300k Supervised Model	3.32	2.69	2.38
300k RL Refined Model	1.66	1.54	1.50
300k SWOR RL Refined Model	1.65	1.53	1.49
Unsupervised Model	1.51	1.48	1.47

Table 4: Empirical comparison of supervised model pretrained on 30k, 150k and 300k programs, their fine-tuned models and our model on 2D CAD dataset with Chamfer distance

model was trained directly on the CAD dataset without supervision with learning rate of 0.01 and entropy coefficient (ec) of 0.009. Entropy coefficient trades off between exploitation and exploration. Higher ec leads to slower convergence but the model is less likely to be stuck at local optimum. Setting ec in the range between 0.005 and 0.012 for this problem has not impact the result significantly. We report the result of the three pretrained models, vanilla RL fine-tuned models, SWOR RL fine-tuned models as well as our model after beam search with $k = 1, 3, 5$ in Table 4.

The poor performance of the pretrained model shows that it is not able to directly generalize to the novel dataset due to the mismatch of the training dataset and the CAD dataset. The unsupervised method exceeds the performance of a supervised model by a large margin because it *was trained directly on the target domain*. It also removes the hyperparameter tuning step in transition to the RL fine-tuning and reaches a result competitive to a refined model. The unsupervised method performs better when $k = 1, 3$ than the fine-tuned models. At convergence, the results of the fine-tuned model pretrained on 300k synthetic data and the unsupervised model become very close at $k = 5$. The two types of refined models reach similar result given same amount of pretraining synthetic data – poorer performance on less data – confirming that the quality of the pretraining




Figure 6: The four examples on the right are from the test set, and the rest on the left are from the training set. The target images are on top and the reconstruction from the output programs are at the bottom.




Figure 7: Compare the entropy estimation following the Equation 4 as a weighted sum of stepwise entropy versus taking the average of the sequence log probability. The number of samples varies from 2 to 80 on the x -axis. The shaded area represents the standard deviation of each estimator. From left to right, we demonstrate the result on datasets of three program lengths.

dataset is a limiting factor for the downstream models.

4.3 VARIANCE STUDY OF ENTROPY ESTIMATION

This study (Figure 7) investigates the empirical relationship between the two variance estimators and verifies Lemma 3.1 that $\hat{\mathcal{H}}_D$ achieves lower variance than $\hat{\mathcal{H}}$ (Section 3.2).

We take a single model saved at epoch 40 during the training time of the length 5, 7, and 9 dataset and estimate the entropy with $\hat{\mathcal{H}}_D$ (Equation 9) and $\hat{\mathcal{H}}$ (Equation 8). We consider two sampling schemes: with and without replacement. We combine both entropy estimation methods with the two sampling schemes creating four instances for comparisons. The x -axis of the plot documents the number of samples to obtain a single estimation of the entropy. We further repeat the estimation 100 times to get the mean and variance. The means of SWR method act as a baseline for the means of the SWOR while we compare the standard deviations (the shaded area) of the two entropy estimation methods.

Across all three datasets, $\hat{\mathcal{H}}_D$ (green) shows significantly smaller variance with the number of samples ranges from 2 to 80. But we notice that longer programs, or more complex images, require much more samples to reduce variance. This makes sense because the search space increases exponentially with the program length. The initial bias in the SWOR estimation dissipates after the number of samples grows over 10 and is greater in dataset with longer program length.

5 DISCUSSION

Current program synthesis approaches for non-differentiable renderers employ a two-step scheme which requires the user to first generate a synthetic dataset for pretraining and then use RL fine-tuning with target images. A purely RL-driven approach does not require the curation of a pretraining dataset and can learn directly from target images. Further, unlike approaches that rely on supervised pretraining, the RL approach does not restrict the model to only the program(s) in the synthetic dataset when multiple equal reconstructions exist. This limitation of the pretraining dataset can further impact the downstream models’ ability to generalize. In this paper we introduced the first unsupervised algorithm capable of parsing CSG images—created by a non-differentiable renderer—into CFG programs without pretraining. We do so by combining three key ingredients to improve the sample efficiency of a REINFORCE-based algorithm: (i) a grammar-encoded tree LSTM to constrain the search space; (ii) entropy regularization for trading off exploration and exploitation; (iii) sampling without replacement from the CFG syntax tree for better estimation. Our ablation study emphasizes the qualitative and quantitative contributions of each design component. Even though our RL approach does not have access to a pretraining dataset it achieves stronger performance than a supervised method on the synthetic 2D CSG dataset. It further outperforms the supervised pretrained model on the novel 2D CAD dataset and is competitive to the RL fine-tuned model.

References

- David Alvarez-Melis and Tommi S Jaakkola. Tree-structured decoding with doubly-recurrent neural networks. 2016.
- Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hind-sight experience replay.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, and Daniel Tarlow. Deepcoder: Learning to write programs. In *International Conference on Representation Learning (ICLR)*, 2017.
- Tony Beltramelli. pix2code: Generating code from a graphical user interface screenshot. In *Proceedings of the ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, pages 1–6, 2018.
- Marc Brockschmidt, Miltiadis Allamanis, Alexander L Gaunt, and Oleksandr Polozov. Generative code modeling with graphs. *arXiv preprint arXiv:1805.08490*, 2018.
- Rudy Bunel, Matthew Hausknecht, Jacob Devlin, Rishabh Singh, and Pushmeet Kohli. Leveraging grammar and reinforcement learning for neural program synthesis. *arXiv preprint arXiv:1805.04276*, 2018.
- Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, 2016.
- Xinyun Chen, Chang Liu, and Dawn Song. Execution-guided neural program synthesis. 2018a.
- Xinyun Chen, Chang Liu, and Dawn Song. Tree-to-tree neural networks for program translation. In *Advances in neural information processing systems*, pages 2547–2557, 2018b.
- Jacob Devlin, Jonathan Uesato, Surya Bhupatiraju, Rishabh Singh, Abdel-rahman Mohamed, and Pushmeet Kohli. Robustfill: Neural program learning under noisy i/o. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 990–998. JMLR. org, 2017.
- Tao Du, Jeevana Priya Inala, Yewen Pu, Andrew Spielberg, Adriana Schulz, Daniela Rus, Armando Solar-Lezama, and Wojciech Matusik. Inversecsg: Automatic conversion of 3d models to csg trees. *ACM Transactions on Graphics (TOG)*, 37(6):1–16, 2018.
- 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*, pages 6059–6068, 2018.
- Kevin Ellis, Maxwell Nye, Yewen Pu, Felix Sosa, Josh Tenenbaum, and Armando Solar-Lezama. Write, execute, assess: Program synthesis with a repl, 2019.
- Yaroslav Ganin, Tejas Kulkarni, Igor Babuschkin, SM Eslami, and Oriol Vinyals. Synthesizing programs for images using reinforced adversarial learning. *arXiv preprint arXiv:1804.01118*, 2018.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- David Ha and Douglas Eck. A Neural Representation of Sketch Drawings. *ArXiv e-prints*, April 2017.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017.
- Zhewei Huang, Wen Heng, and Shuchang Zhou. Learning to paint with model-based deep reinforcement learning. *arXiv preprint arXiv:1903.04411*, 2019.
- Philip M Hubbard. *Constructive solid geometry for triangulated polyhedra*. 1990.
- Kacper Kania, Maciej Zięba, and Tomasz Kajdanowicz. Ucsq-net—unsupervised discovering of constructive solid geometry tree. *arXiv preprint arXiv:2006.09102*, 2020.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. *arXiv preprint arXiv:1802.05983*, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Wouter Kool, Herke Van Hoof, and Max Welling. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International Conference on Machine Learning*, pages 3499–3508, 2019.
- Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1945–1954. JMLR. org, 2017.
- Adam Daniel Laud. Theory and application of reward shaping in reinforcement learning. Technical report, 2004.

- Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd gan: Towards deeper understanding of moment matching network. In *Advances in Neural Information Processing Systems*, pages 2203–2213, 2017.
- Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. Differentiable monte carlo ray tracing through edge sampling. In *SIGGRAPH Asia 2018 Technical Papers*, page 222. ACM, 2018.
- Hsueh-Ti Derek Liu, Michael Tao, Chun-Liang Li, Derek Nowrouzezahrai, and Alec Jacobson. Beyond pixel norm-balls: Parametric adversaries using an analytically differentiable renderer. 2019.
- Yunchao Liu, Zheng Wu, Daniel Ritchie, William T Freeman, Joshua B Tenenbaum, and Jiajun Wu. Learning to describe scenes with programs. 2018.
- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. *arXiv preprint arXiv:1811.12359*, 2018.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016.
- Emilio Parisotto, Abdel-rahman Mohamed, Rishabh Singh, Lihong Li, Dengyong Zhou, and Pushmeet Kohli. Neuro-symbolic program synthesis. *arXiv preprint arXiv:1611.01855*, 2016.
- Lukasz Romaszko, Christopher KI Williams, Pol Moreno, and Pushmeet Kohli. Vision-as-inverse-graphics: Obtaining a rich 3d explanation of a scene from a single image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 851–859, 2017.
- Gopal Sharma, Rishabh Goyal, Difan Liu, Evangelos Kalogerakis, and Subhransu Maji. Csgnet: Neural shape parser for constructive solid geometry. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5523, 2018.
- Eui Chul Shin, Illia Polosukhin, and Dawn Song. Improving neural program synthesis with inferred execution traces. In *Advances in Neural Information Processing Systems*, pages 8917–8926, 2018.
- Richard Shin, Neel Kant, Kavi Gupta, Christopher Bender, Brandon Trabucco, Rishabh Singh, and Dawn Song. Synthetic datasets for neural program synthesis. *arXiv preprint arXiv:1912.12345*, 2019.
- 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. *arXiv preprint arXiv:1901.02875*, 2019.
- Shubham Tulsiani, Hao Su, Leonidas J Guibas, Alexei A Efros, and Jitendra Malik. Learning shape abstractions by assembling volumetric primitives. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2635–2643, 2017.
- Aaron van den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pages 6306–6315, 2017.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Jiajun Wu, Joshua B Tenenbaum, and Pushmeet Kohli. Neural scene de-rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 699–707, 2017.
- Shunyu Yao, Tzu Ming Hsu, Jun-Yan Zhu, Jiajun Wu, Antonio Torralba, Bill Freeman, and Josh Tenenbaum. 3d-aware scene manipulation via inverse graphics. In *Advances in Neural Information Processing Systems*, pages 1887–1898, 2018.
- Pengcheng Yin and Graham Neubig. A syntactic neural model for general-purpose code generation. *arXiv preprint arXiv:1704.01696*, 2017.
- Amit Zohar and Lior Wolf. Automatic program synthesis of long programs with a learned garbage collector. In *Advances in Neural Information Processing Systems*, pages 2094–2103, 2018.