

Cover Letter for PSML 2020 Project Report:

Dear Instructors,

Our PSML project discussions began with this idea: **Synthesizing Sketches from Logical Formulae using Machine Learning techniques**. More specifically, we were investigating SyGuS (Syntax Guided Synthesis) benchmarks for logical specifications. These discussions led to two interesting streams of thought: one was a data driven approach using Deep Learning whereas the other was a policy learning approach using Reinforcement learning techniques. During the course of our discussions, we decided to go ahead with the former for the PSML project. Meanwhile, Ravi simultaneously decided to go ahead with a mini literature survey of the latter branch for his Machine Learning course project term report (of atmost 4 pages) where the requirement was to survey existing work. Since the latter stemmed from PSML course discussions, we decided to append that report, for completeness, in this document as well.

Consequently, please find enclosed the following two documents:

- PSML Report: Sketches for Logical Specifications using Deep Learning
- Learning Program Sketches from Logical Specifications using Reinforcement Learning

- Ravi Raja and Stanly Samuel (22 June 2020)

PSML Report: Sketches for Logical Specifications using Deep Learning

Chiranjib Bhattacharyya, Deepak D’Souza, Ravi Raja, Sriram Rajamani, and Stanly Samuel

Abstract—We consider the synthesis of sketches for logical formulae using Deep Learning techniques. For this problem in program synthesis, we discuss relevant literature, gaps in existing techniques and our approach to attack the problem. This report broadly consists of a research proposal and preliminary progress in this direction. For the preliminary direction, we discuss a prototype tool as a work in progress.

I. INTRODUCTION

Recent years have seen a surge of interest in the area of program synthesis where generating correct-by-construction programs with respect to high-level specifications is the end goal. In the current decade, these techniques have been proven beneficial to end users which was earlier not possible, especially during the nascent years of this area in the 1950’s. Recently, there has been a lot of work in the intersection of Machine Learning with Program Synthesis and this is where this project lies.

Although there are many techniques in this area, we will look into the sub-area of programming by sketches which is a Good Old Fashioned Program Synthesis (GOFPS) technique and its automation using Deep Learning techniques.

Sketches are partial programs that aid the process of program synthesis from specifications. They restrict the search space of programs for the synthesizer thereby speeding up the synthesis process. However, the speed-up is contingent on the fact that we synthesize the right sketch for the input specification. Intuitively, a right sketch is a partial program for which there exists at least one solution i.e a complete program that satisfies the specification for all inputs, obtained by completing the sketch.

Sketches were first introduced in [17]. In this work, a user specifies a sketch which is a partial program. and a specification that constrains the behaviour of the sketch. The sketch synthesizer then converts the sketch and specification into a boolean formula which is fed into a SAT solver. This work is the first to describe the Counter Example Guided Inductive Synthesis (CEGIS)

architecture and is used to synthesize holes that are constants. Here, the user is expected to provide the Sketches. Providing sketches as well as the specification may be inefficient as the user may perform redundant work. It is thus beneficial to at least generate sketches automatically when the specification is given.

SketchAdapt [14] generates sketches automatically using statistical techniques. This work is inspired from RobustFill [9] and Deepcoder [3]. The key idea here is to generate program sketches from user specifications i.e. I/O examples or natural language that flexibly manage the work load between neural synthesis and symbolic search. Sketches are valid program trees in the DSL, where any number of sub-trees has been replaced by a special token called *HOLE*. Intuitively, this token designates locations in the program tree for which pattern-based recognition is difficult, and more explicit search methods are necessary. This tool can be divided into two components: 1) *Neural Sketch Generator* which is a Seq2Seq neural architecture. It takes the spec. as input sequence and produces a good sketch as a sequence of tokens, and 2) *Program Synthesizer* which takes the generated sketch and outputs a final program by choosing grammar productions based on learned production probabilities. Generating a good sketch prunes out a large part of program space, thereby improving the synthesis time. Due to use of symbolic search to fill up the sketches, programs produced are more generic. While this work is a huge success for the domains of List Processing and String Processing, it does not work for logical formulae. Further, it requires millions of training data for training the network which is not readily available for logical domains. Moreover, I/O specifications are finite and do not capture an infinite state space. Thus, synthesizing sketches for logical specifications remain an open problem.

The closest approach for synthesizing sketches for logical specifications is that of Counter Example Guided Inductive Synthesis modulo Theories i.e CEGIS(T) [1]. In this work, the solver invokes CEGIS to synthesis a sketch or a program skeleton as they call it. A sketch in this case is a program synthesized by an enumerative solver where the constants are replaced with holes. The sketch is then fed into a constraint solver to check feasibility. If feasible, the constraint solver figures out the right

C. Bhattacharyya, D. D’Souza, R. Raja and S. Samuel are with IISC, Bengaluru, India.

S. Rajamani is with Microsoft Research, Bengaluru, India.

constants and generates the final program. Otherwise, the sketch will not have any valid completion that satisfy the logical specification. In this case, the "bad" sketch is appended as additional information to the original constraint. Their tool FastSynth implements this approach. However, this approach can only synthesize sketches from I/O examples which are sampled from the logical specification. Moreover, the sketches obtained are not really sketches because they are whole programs replaced by holes. Also, the synthesizer for CEGIS(T) does not really encode the complete semantics of the logical specifications.

Counter Example Guided Neural Synthesis [15] extends the CEGIS loop to synthesize programs in the counterexample loop using a neural network. Even this case the neural network uses I/O example embeddings and not Logical Specifications.

Another work which is close to our idea is the Multi-modal Synthesis of Regular Expressions [5] where they synthesize regular expressions from two modes of specifications i.e. natural language and I/O examples respectively. They first use a semantic parser to generate from English descriptions a type of sketch which they call a hierarchical sketch (say S). Given S , their PBE engine synthesizes a regular expression that is a valid completion of S while satisfying the I/O specification. Although this work shows promise, it is likely that the user intent may be misrepresented due to the fact that natural language can be inherently ambiguous. Moreover, these specification do not capture an infinite space of examples. Thus, being able to use logical formulae to aid the process of sketch synthesis is of importance.

Based on this exhaustive survey, we propose the following problem statement:

Problem Statement: Given a logical specification and a DSL for the space of programs, we wish to design a system that synthesizes sketches in such a way that the work load between neural synthesis and symbolic search is managed efficiently. We wish to design a neural network for the same and use embeddings of logical formulae. To the best of our knowledge, we are the first to work on sketch generation using logical embeddings.

Other works using sketches in some form: Neural Edit Completion [?] is a code completion technique that uses the context of the code environment as specification. The Sketch is the partial code Program synthesis using conflict-driven learning [10] also talks about sketches in terms of partial programs.

Contribution: In this report, we will talk about a research proposal and preliminary work in this direction. The major topics are:

- A line of direction to generate synthetic dataset for SyGuS benchmarks.
- A "concolic" neural network architecture using logical embeddings and I/O embeddings.
- A prototype tool SyGuS-Sketcher using existing techniques with experimentation plan.

II. MOTIVATION

One of the motivating examples that we had looked into involved the Conditional Linear Integer Arithmetic track of SyGuS benchmarks. A state of the art solver CVC4 was not able to solve the following constraint: $x \text{ is even} \implies \text{div}(x) + \text{div}(x) = x$ where $\text{div} : \text{Real} \rightarrow \text{Real}$ is the function to be synthesized. The constraint hints on the fact that the division operator needs to be used in the final program. A human would have had no issue in figuring this out. Consequently, we expect to design and train a neural network that could pick up these patterns and synthesize sketches accordingly. For example, feeding this constraint to the neural network could synthesize a sketch as follows: $\text{div}(x) = x/\text{HOLE}$ i.e. the neural network realizes that the division operator is to be used but doesn't really know which constant is needed so it synthesizes a hole in place of the constant. A traditional sketch solver can then be used to synthesize the constant 2.

On the other hand, the neural network may even choose to synthesize the sketch $\text{div}(x) = \text{HOLE}/\text{HOLE}$ in which case you will also need an enumerative solver to fill up the hole for the variable x .

The other motivation is to improve the performance of existing tools by the use of sketches. In the experiments section, we will comment on DeepSynth's performance on its benchmarks.

III. BACKGROUND

In this section, we explain the background required for our report. We explain CEGIS [17], CEGIS(T)[1] and CEGNS[15]. We also touch upon SyGuS [2] benchmarks in a nutshell.

For brevity, we refer the readers to the above papers.

IV. RESEARCH PROPOSAL

In this section, we discuss our research idea in detail. We first discuss the overall idea of our research using a block diagram. We then discuss a neural network architecture that we propose for the problem. We also explain DeepSynth's NN architecture for clarity.. Finally, we conclude with a discussion on how to generate a dataset to train this problem. The benchmarks that we are looking into are the standard SyGuS benchmarks.

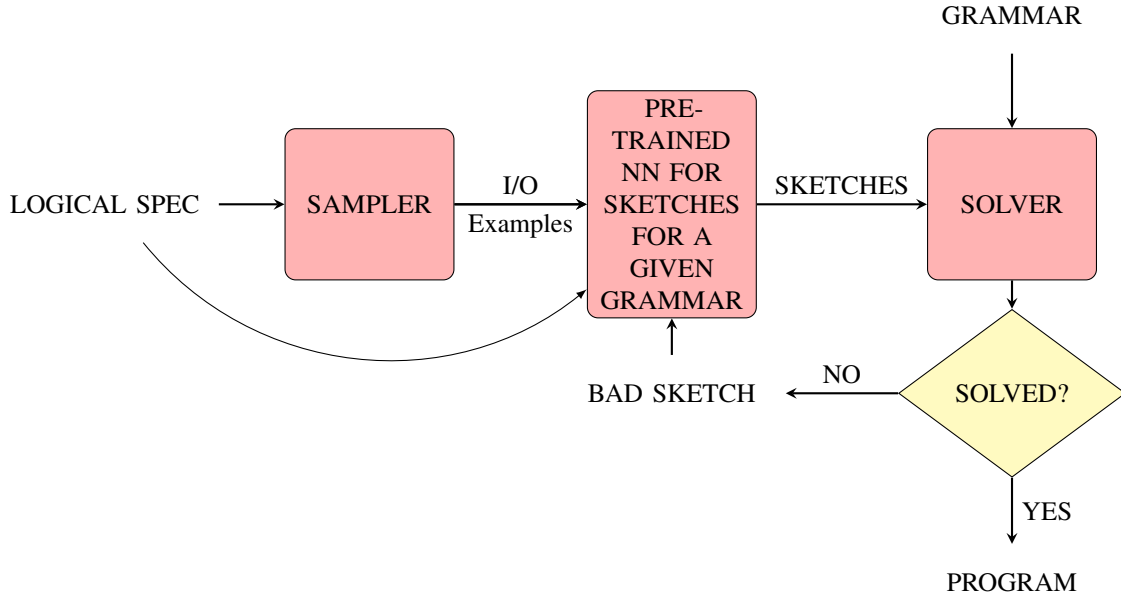


Fig. 1. Synthesizing Sketches from Logical Embeddings

A. Overall idea

Figure 1 describes the high level idea for synthesizing sketches from logical embeddings. We first sample a finite set of I/O examples from the logical specification. We have a pre trained neural network for a given DSL that is trained to generate the required sketches when a logical specification and I/O specification are given as input. As a novel extension, it may also be possible for the NN to take as input a bad sketch and make decisions accordingly. This would be something similar to CEGIS(T) but using an NN. We then feed it to a traditional solver that solves the sketch using enumerative techniques if the sketch has non constant holes and constraint based techniques, if the sketch has constant holes. If the sketch is infeasible, then we use this to direct our NN to synthesize a better sketch. The meaning of an infeasible sketch is that there does not exist a valid (w.r.t. the DSL) completion of the sketch for the given specification.

The major research question here is how to make the NN understand whether a given sketch is an infeasible sketch or not.

B. Proposed Neural Network Architecture

This section details about the proposed Neural Network architecture for our work 2. Sygus-Sketcher uses Neural Network for two purposes - 1) Learning representation, and 2) Synthesizing Sketches from the learned representations. These are elucidated below:

1) *Neural Network Architectures:* This section explains different architectures to be used in this work.

a) *Gated Graph Neural Network (GGNN):* We use Gated Graph Neural Network (Li et al., 2015)[13] for representing specifications. This part gives a brief understanding of how GGNN works. Let $G = (V, E, X)$ be a graph with V being set of vertices, $E = (E_1, E_2, \dots, E_k)$ be the set of directed edges where k is the number of edge types, and X is the feature set of nodes. In our case $k=2$ types of edges viz. intrinsic and extrinsic depending upon whether it is connecting nodes internal to grammar or logical spec or is it connecting the grammar and logical spec graphs together. Each vertex $v \in V$ is annotated with feature vector $x^{(v)} \in R^d$ representing the features of node (e.g. embedding of node label).

Each of the node is associated with a state embedding vector $h^{t(v)}$ at time step t , initialized with node feature vector $x^{(v)}$. The sizes of State vector and feature vector are kept same (R^d). Information is propagated among each nodes of the graph by Message Passing. Each node v receives messages of type k from its neighbours u , where each message is computed from its current state vector as $m_k^{t(u)} = f_k(h^{t(u)})$. Here, f_k can be an arbitrary function which we choose to be a Linear Layer that transforms the state vector with a $d \times d$ weight matrix W as $f_k(h^{t(u)}) = (W \cdot h^{t(u)})$. All the state vectors are updated at same time by aggregating all the messages from its neighbours as $M^{t(v)} = g(m_k^{t(u)})$ at time step

t, where g is an aggregation function which is element-wise summation in our case. A new state vector $h^{t+1(v)}$ for each vertex v at next time step t+1 is calculated as $h^{t+1(v)} = \text{GRU}(M^{t(v)}, h^{t(v)})$, GRU is Gated Recurrent Unit (Cho et al.) [8]. These updates are carried out for a fixed time step T and the final state vectors are used as node representation. To get a graph level representation \mathcal{G} we do max pooling of final node representations as $\mathcal{G} = \text{MAXPOOL}(h^{t(v_i)})$ for each node v_i .

b) *Gated Recurrent Unit*: Gated Recurrent Unit or GRU cell is an LSTM cell (explained below) with gates for controlling the amount of information to be stored in, forget from the cell state (memory of LSTM cell) of the LSTM cell.

c) *Seq2Seq Architecture*: It is an Encoder-Decoder model, where the encoder encodes the input sequence (I/O Examples in our case) into an embedding vector which is passed as input to the Decoder network. The Decoder then decodes the output sequence (Sketch in our case). We use Long Short Term Memory (LSTM) both for Encoding and Decoding.

d) *RNN*: Recurrent Neural Networks or RNN are generally applied to sequential data. Given a Sequence of input (x_1, \dots, x_T) , an RNN computes the sequence of outputs (y_1, \dots, y_T) by iterating the following equation:

$$h_t = \text{sigmoid}(W^{hx}x_t, W^{hh}h_{t-1})$$

$$y_t = W_{h_t}^{yh}$$

where h_t represents hidden state of RNN.

The RNN can easily encode sequences to sequences whenever the alignment between the inputs the outputs is known ahead of time (i.e both having same lengths). However, the application of an RNN to problems whose input and the output sequences have different lengths is not very clear. Cho et al. [7] uses one RNN to encode input to a fixed-size vector and uses another RNN to generate output sequence from the fixed-size vector. Theritecally this approach should work but it is seen practically that it fails to recall long-term dependencies ([4], [11]). However, LSTM's [12] are able to capture the long term temporal dependency in the sequence.

e) *LSTM*: LSTM estimates the conditional probability $p(y_1, \dots, y_{T'} | x_1, \dots, x_T)$ where (x_1, \dots, x_T) is an input sequence and $y_1, \dots, y_{T'}$ is an output sequence whose length T' may be different from the length of input sequence. The LSTM computes this conditional probability by first by learning a fixed-dimensional representation v of the input sequence (x_1, \dots, x_T) which is the last hidden state of the LSTM. This representation vector v is then passed to another LSTM, as an initial hidden state, that serves as a Language Model (trained

over the Grammar of required output language) to compute the probability of $y_1, \dots, y_{T'}$

$$p(y_1, \dots, y_{T'} | x_1, \dots, x_T) = \prod_{t=1}^{T'} p(y_t | v, y_1, \dots, y_{t-1})$$

Where $p(y_t | v, y_1, \dots, y_{t-1})$ distribution is represented with a softmax operation over the vocabulary (unique tokens in the output grammar G). The first LSTM is the Encoder part of a Seq2Seq architecture while the second one is the Decoder part.

2) *Learning Representations*: The Constraints or Semantic/Logical Specification and the Grammar or Syntactic Specification are highly structured, which is missed when we consider them just a sequence of tokens. The structure contains important information like order of operators in the constraints or derivations of the grammar. However, it may fail to see the semantic similarity of syntactically different constraints. Consider the following three constraints -

- (1) $(\implies (= (\text{mod } x \ 2) \ 0) (= (+ (\text{div1 } x) (\text{div1 } x)) x))$
- (2) $(\text{not } (\implies (= (\text{mod } x \ 2) \ 0) (= (+ (\text{div1 } x) (\text{div1 } x)) x)))$
- (3) $(\implies (\text{True}) (\implies (= (\text{mod } x \ 2) \ 0) (= (+ (\text{div1 } x) (\text{div1 } x)) x)))$

Assume same grammar G for all three constraints. Constraints are said to deliver the semantics of the function to be synthesized but still due to its symbolic nature, constraints (1) and (2) are considered syntactically much more similar than constraints (1) and (3) despite the fact that (1) and (3) are essentially asking for same function i.e division by 2.

While I/O examples can encode concrete functional behaviour of intended function to be synthesized. For e.g. (4, 2), (10, 5) I/O pairs clearly concludes that (1) and (3) are semantically equivalent. But I/O examples suffers from insufficient path coverage. And even for covered paths, it requires huge dataset to achieve generalization, leading to expensive training procedure.

Therefore we aim to benefit from both kinds of specifications viz. SMT and I/O examples. For this we present a novel approach for learning representations of both SMT spec and I/O spec. This idea is based on Learning Blended, Precise Semantic Program Embeddings by Wang et al. [18]

a) *Encoder: Learning Representations for Logical spec and given Grammar*: In order to retain the structural properties of constraints and grammar we construct GGNN for the underlying AST's. Each node of the GGNN is a GRU cell and each edge is a feed forward network. It is trained and the final graph embedding \mathcal{G} are obtained as explained in IV-B.1.a. Our architecture

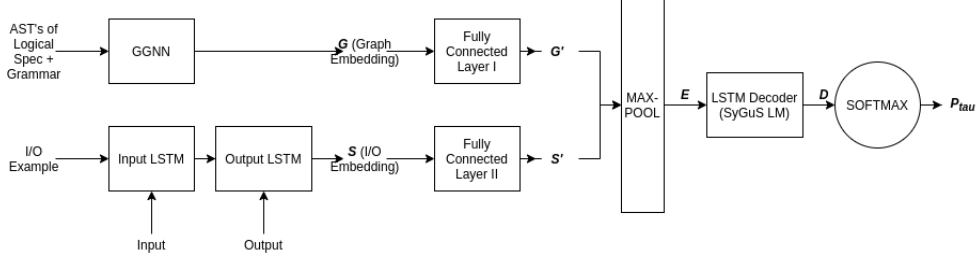


Fig. 2. Proposed Neural Network Architecture. The dimensions of all the hidden state vectors \mathcal{G} , \mathcal{S} , \mathcal{G}' , \mathcal{S}' , \mathcal{E} , \mathcal{D} , is 128-by-1 and size of P_τ is equal to the size of Vocabulary. The Encoder gives output as \mathcal{E} , which is then passed to the Decoder LSTM to get the program token probabilities P_τ .

for learning representations for syntactic and semantic specification is similar to the one used by Xujie Si [16] which joins the graphs of logical specification and the grammar together as one graph. In addition, we extend it for General Track of SyGuS Competition.

b) *Encoder: Learning Representations for I/O spec:* We use an **input LSTM encoder** and an **output LSTM encoder** as described in deepsynth [15]. Each argument of the input(I) of an I/O example is fed to the input LSTM. The final hidden state of the input LSTM encoder becomes the initial hidden state of the output LSTM encoder and the output (O) is fed as input to it. The final hidden state of the output LSTM is the learned representation of one I/O example. Embeddings for all the I/O examples of a program are aggregated together to get the final I/O embedding \mathcal{S} .

c) *Aggregating \mathcal{G} and \mathcal{S} :* Finally we learn a combined representation for both logical spec and the I/O spec by aggregating the two embeddings \mathcal{G} and \mathcal{S} by first passing them through a fully connected neural network and then performing max pooling operation on them. We denote this final embedding by \mathcal{E} .

3) *Decoder: Neural Network for Synthesizing sketches:* Final embedding \mathcal{E} from the Encoder is then passed to an LSTM decoder, which is a Language Model trained over SyGuS Grammar, to produce final hidden representation \mathcal{D} . We pass this embedding \mathcal{D} through a Softmax layer to get the token probability distribution P_τ where τ denotes the set of tokens. The Decoder at each time step t predicts the next most probable token τ_i . τ_i is given as input to the LSTM decoder for the next time step. At each time step token with highest probability is generated. This is called Greedy Decoding technique. Greedy Decoding is an irreversible process i.e. if a wrong token is generated then it will lead to a wrong sketch, there's no way to backtrack. Therefore, we use Beam Search Decoding, which searches for \mathcal{K} most probable tokens at each time step and generates \mathcal{K} most likely sketches.

4) *Solving Sketches and Training the Neural Network:* The sketches generated by the Decoder is given to a Solver which finally generates \mathcal{K} complete programs $P_{\mathcal{K}}$. These are then compared with the true program \mathcal{P} and cross-entropy loss is calculated as follows:

$$\mathcal{L} = - \sum_{i=1}^d \mathcal{P}_i \cdot \log(\mathcal{P}_{\mathcal{K}_i})$$

where $P_{\mathcal{K}}$ and \mathcal{P} are 128 dimensional vectors denoting the output programs and the ground truth programs. \mathcal{P}_i and $\mathcal{P}_{\mathcal{K}_i}$ are the i^{th} dimension value of \mathcal{P} and $P_{\mathcal{K}}$ respectively. We train our network by the method of *teacher forcing*, in which the target sequence is fed to the Decoder directly during the training and used to compute the next token. The final program P_{final} obtained by minimizing the loss \mathcal{L} , is our solution program.

C. DeepSynth

This section explains the neural network architecture and training data generation techniques used in deepsynth and how are these integrated together.

1) *Neural Network:* The DeepSynth Neural Network, shown in Fig. 3 is made up of several parallelised sequence-to-sequence neural networks, each of which processes a single counterexample (i.e. I/O example) and computes a probability distribution over the likely sequence of tokens in the program given that input/output example. The outputs from all the sequence-to-sequence networks are combined using pooling. In the following sections we describe in detail a single sequence-to-sequence network, and then the pooling process.

a) *Processing a Single counterexample:* The Seq2Seq network used for processing a single counterexample is shown in Fig. 4 and consists of:

- 1) An **input encoder** LSTM network, which receives the input argument values from the I/O example.
- 2) An **output encoder** LSTM cell, which receives the output value from the I/O example and is

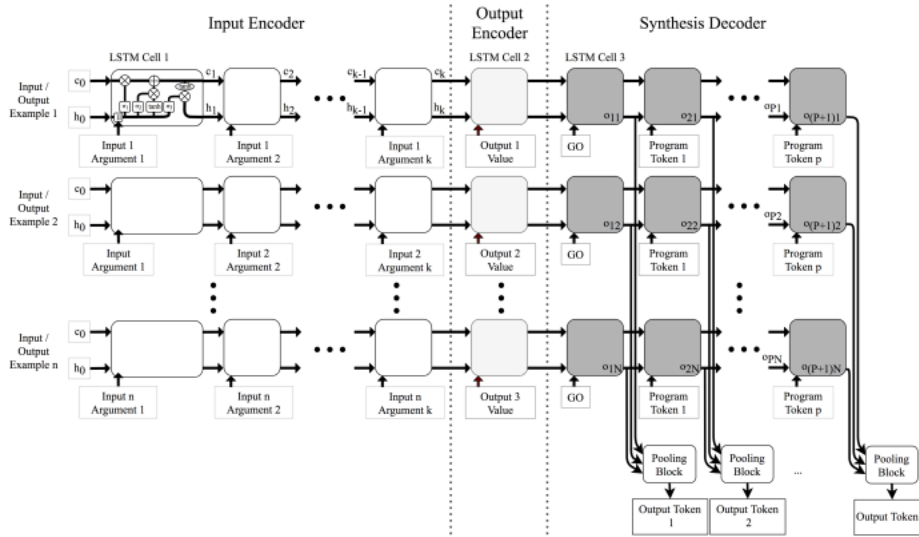


Fig. 3. The DeepSynth Neural Network

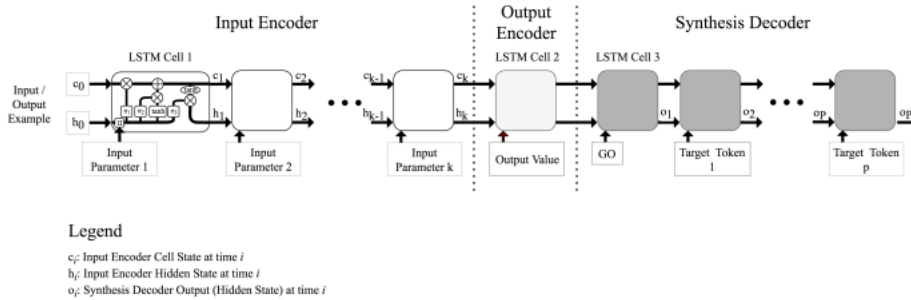


Fig. 4. Seq2Seq network for a single I/O example

conditioned on the final cell and hidden states of the input encoder network.

- 3) A **synthesis decoder** LSTM network, which is trained on the **target program** and conditioned on the final output encoder state.

Input arguments are fed into the input encoder network as a sequence. The cell state and hidden states of all three LSTM cells (for the input encoder, output encoder, and synthesis decoder respectively) are set to be 128-dimensional vectors. Program tokens are represented using a learnable 128-dimensional vector embedding. Every program token in the token vocabulary has a 128-dimensional vector representation which is updated at each time step of training.

DeepSynth evaluates three different representations for input arguments and the output values, which is referred

as **input-modes**:

- 1) **"normalised"** mode, where every parameter value p is represented by a normalised scalar value, more specifically $\frac{p}{2^{31}} - 1$
- 2) **"binary"** mode, where p is provided in its binary representation, a 32-dimensional vector of 0s and 1s. This representation is more useful for helping the network detect minute differences between inputs.
- 3) **"normBinary"** mode, which combines both normalised and binary representations, thus producing a 33-dimensional vector representation for every input argument and the output value.

Finally, it includes an additional setting for the network, in which the synthesis decoder is provided with the arity (i.e. number of input arguments) of the target program in addition to the target program's tokens.

b) *Pooling*: To process n I/O examples, the DeepSynth neural network creates n identical copies of the single-IO Seq2Seq network, with each receiving one of the n I/O examples. Once it has the encoded I/O examples, it is then fed to n synthesis decoders along with a GO token at time step 1, and a 128-dimensional hidden state is produced from each decoder. Let o_{t_i} be the hidden state produced at time t by synthesis decoder i .

To aggregate all hidden states o_{T_i} at any given time T , these values are first passed through a **fully-connected neural network layer** to compute n values fc_{T_i} such that $fc_{T_i} = \tanh(W * o_{T_i} + b)$, where W is a 128-by-128 weight matrix and b is a 128-by-1 bias matrix. Following this, the n fc_{T_i} values are aggregated into a unique 128-dimensional aggregate output value O_T using the max pooling operation as follows:

$$O_T[j] = \max_{1 \leq i \leq n} fc_{T_i}[j]$$

where $O_T[j]$ is the j^{th} dimension of O_T and $fc_{T_i}[j]$ is the j^{th} dimension of fc_{T_i} . Finally, O_T is used to compute the probability distribution D_T over the next token using a softmax neural network layer. More formally, $D_T = \text{softmax}(W_{Lin} * O_T + b_{Lin})$, where W_{Lin} is a 50-by-128 weight matrix and b_{Lin} is a 50-by-1 bias vector (50 is the size of vocabulary). The pooling procedure is shown in Fig. 5

2) *Training Data Generation*: The Training Data consists of a set of Candidate programs, each accompanied by a corresponding set of input/output examples. For program synthesis, neural networks typically require millions of training examples therefore, in order to obtain a sufficiently large training data set, they designed a program generator that randomly generated programs made up of syntactically correct combinations of SYGUS-IF instructions¹. Randomly generated data set also consists of **”bad training data”** which is then pruned out using some set of rules and SMT based procedure. **Bad Training Data**: Redundant programs and input/output examples that do not sufficiently differentiate between programs. These are removed using following methods.

a) *Eliminating redundancy in programs*: **Constants in programs** There are several cases where the choice of constants in programs can introduce redundancy.

- Shifting a bit vector by any number greater than the bit vector width produces a 0. Thus, a rule is introduced such that no generated program contains any shift operation with the second operand greater

than the width of the first.

- Zero value constants are disallowed in all cases since they almost always introduce redundancy.

SMT based Redundancy Check SMT-solver is used to identify two sources of redundancy:

- An if-then-else statement is redundant when it always returns one of its conditional outputs. This redundancy is identified by recursively identifying leaf operators (i.e. operators not having nested ife statements) and checking their branch satisfiability with an SMT solver. If a branch is unreachable, the if-then-else statement is replaced with the other branch and call the recursive function again. This is continued until no leaf statements are redundant.
- The SMT Solver also identifies those programs that always return a single constant value, or always return a single one of the input arguments, and remove them from the training set.

b) *Input/Output Example Generation*: Input/Output examples are generated uniformly from a target program by executing the programs in order to get the corresponding outputs.

DeepSynth generates Input/Output (I/O) examples for each random program such that this I/O is informative with respect to this program, i.e., the I/O must cover all conditional branches in a program so as to fully portray a program’s semantics. This is done by using a combination of randomly generated inputs and SMT solving. For a program P with C conditional statements and a shift assertion set S , “Smart” mode cycles through all 2^C possible execution cases and, where a case is satisfiable subject to S , it produces a new I/O example. This is done until N I/O examples are produced. In order to obtain a distribution of inputs over the input space when using an SMT solver, Z3 is used with the phase selection set to random.

c) *Program Tokenisation*: Programs are converted into a **token** representation that the network can process. This representation encodes every operator in the DSL, as well as up to 10 different input parameters in a program, as its own token. Constants are each encoded using 8 tokens, such that each token represents a 4 bit value and can take on one of 16 values. Two extra tokens are introduced for GO and EOS, which is used by the network during training to learn when synthesis starts/stops. To uniformise the length of its program sets, a PAD token is introduced. In total the vocabulary consists of 50 tokens for representing operations, input parameters, and constants that can appear in the program DSL.

d) *Batches*: Programs’ token representations and I/O examples are aggregated into batches of size B .

¹<https://sygus.org/assets/pdf/SyGuS-IF.2.0.pdf>




Fig. 5. The Pooling Operation

It also pads batch programs to one same length for computational reasons and converts I/O values into normalised and binary input format.

D. Generating training data

Our aim is not to replace IO with generalized specifications but to aid these specifications with I/O examples. Generating training data has been the bottleneck for this research work since it is not a trivial pursuit in our case. Consequently, if we manage to generate synthetic SyGuS training data, it can prove to be a substantial contribution for Deep Learning based SyGuS research. Thus, generating data is one of the ideas we will focus on in this report. As discussed in the previous section, DeepSynth’s Neural Network architecture requires training data which is a pair of finite I/O specification and correct program. They achieve this by first randomly generating programs and inputs with certain restrictions. Next, they generate the outputs by feeding these inputs to the generated programs. Thus, it is relatively easier to generate programs this way. Moreover, they claim to have promising results using this strategy. This is encouraging for us to try a similar strategy for the generation of synthetic sygus benchmarks as well. However, we need training examples of the form (logical constraint + IO examples, correct program). The only option is to generate programs by feeding the constraints to different solvers. A natural question to ask would be how to generate millions of constraints? This is a challenging problem and will require a new data generation pipeline. We will discuss some ideas for data generation in this section. Our aim is to train the Neural Network with constraint-program pairs in such a way that the NN can generalize well. Thus, it is important that we feed various combinations of the existing constraints to the NN.

Mutations of existing SyGuS benchmarks: We can exploit properties of logical formulae such as commutativity, associativity etc to synthesize semantically

equivalent but syntactically different constraints. For example, the formula A and B can be represented as B and A. Such mutations to existing formula will help in generation of constraints. Moreover, it should generate the same program which will aid the training of the NN.

Using different solvers: The other possibility is to use different solvers so that they generate different programs for the same constraint. We have made this observation while experimenting with different tools. Correctness of these programs is important as well and hence a reliable solver such as CVC4 is important. As per our experiments, some solvers such as DryadSynth were not sound.

Bounding synthesis time: Assume we require 80 million training examples as per DeepSynth. In such a case, we need 80 million constraints randomly generated using an intelligent strategy. Assuming we have access to a standard solver that can solve each constraint in approx 10 ms (which is a realistic expectation), the solving should take approx 10 days on a basic machine whereas increasing it to 50 ms may need 52 days of solving and 1 sec bound leads to 2.5 years of solving. Thus, bounding the synthesis time for constraints is important. More importantly, it is important to generate constraints that will be solvable within these bounds.

Random generation of constraints: The previous ideas may not scale to millions of training data that we need. Thus, at some point, we have to look into generating random constraints using ideas from DeepSynth as a starting point. This becomes challenging in our situation as the structure of the constraints change for different classes of SyGuS benchmarks. For example, Loop Invariant SyGuS benchmarks have a PRE, TRANS, POST format whereas the others do not. Thus, it is prudent to decide a class of benchmarks first and proceed.

Number of training examples: Since we are using logical constraint and I/O specification as an input to the

NN instead of just I/O specifications, it may be possible that we require fewer examples for training. However, this can only be empirically evaluated.

V. SYGUS-SKETCHER: A PRELIMINARY PROTOTYPE

Instead of solving the larger problem as mentioned in the previous section, we decided to solve a sub-problem first. We have begun implementing a preliminary prototype as a tool which we call SyGuS-Sketcher. SyGuS-Sketcher is built atop DeepSynth [15] and is available on github.² In this section, we discuss the tool architecture for SyGuS-Sketcher and experimentation direction. The aim of this prototype is to show that sketches for logical specifications can indeed help improve the performance and number of benchmarks solved as compared to DeepSynth. This is our hypothesis.

A. Tool architecture

Figure 6 shows the prototype architecture. As you can see, the architecture is similar to DeepSynth except for two modifications. First, out of the top K candidates selected by the DeepSynth’s NN, we select the candidate most likely to be correct, which is a complete program. Next, we replace the constant values with constant literals to generate a sketch. This is as good as assuming that we do not trust the constants generated by the Neural Network but use existing sketch compilers to fill up the program while guaranteeing the correctness of the specification. Thus, this boils down to a problem of solving sketches with constant holes which is exactly what the Sketch tool[17] does. However, for the bit-vector invariant generation benchmarks that we use, Sketch has limited support. For example, bit-vector subtraction and relational operators are not supported in Sketch due to which we had to look for other options. Thus, we looked into CEGIS(T) which has this support. CEGIS(T). We are currently in the process of implementing this integration. This is the second modification from DeepSynth.

B. Planned experiment

We plan to test this prototype against DeepSynth’s 88 bit-vector invariant benchmarks. We describe a subset of these benchmark timings below:

²<https://github.com/stanlysamuel/sygus-sketcher>

Benchmarks	DeepSynth	SyGuS-Sketcher
anfp-new	23.7529s	.
formula22	44.5536s	.
hola.05	165.086s	.
array-new	76.1318s	.
cegar2-new	106.402s	.

As you can observe, DeepSynth takes a considerable amount of time whereas traditional solvers without neural network architectures have an order in milliseconds. This is where we expect Sketches to play a role in improving the performance time.

VI. FUTURE WORK AND EXTENSIONS

We wish to achieve the following in the coming weeks:

- Complete the implementation of SyGuS-Sketcher and compare it empirically with DeepSynth.
- Generate dataset as explained and train the proposed NN model for invariant benchmarks and perform relevant experimentation.
- Extend to more classes of SyGuS benchmarks such as CLIA to check if this architecture can fill the gaps, as mentioned in the motivating example.

Research directions:

If a sketch generated by the neural network is infeasible (i.e. the solver is not able to find a valid completion of the sketch that satisfies the specification), how do we communicate it back to the neural network? More specifically, how do we leverage a neural network to understand infeasible sketches, given an embedding of a bad sketch using a GGNN? A deduction guided RL approach [6] uses a similar idea but to update the policy of an RL network when an infeasible partial program is returned.

The verifier that SyGuS-Sketcher uses can only synthesize constant holes. Consequently, infeasible sketches can be encoded as logical formulae with constant literals and fed as a constraint back to the synthesizer. However, if the hole is not a constant hole, how do we regard it as a bad sketch is still not clear. One plausible direction is to look into the use of Conflict Driven Learning techniques for the same.

VII. CONCLUSION

In this report, we have proposed a research idea and a plausible line of attack. We have also started preliminary work in implementing the prototype and experimental evaluation is under way.

This report highlights the following contributions: 1) A line of direction to generate synthetic dataset for SyGuS benchmarks, 2) A "concolic" neural network architecture

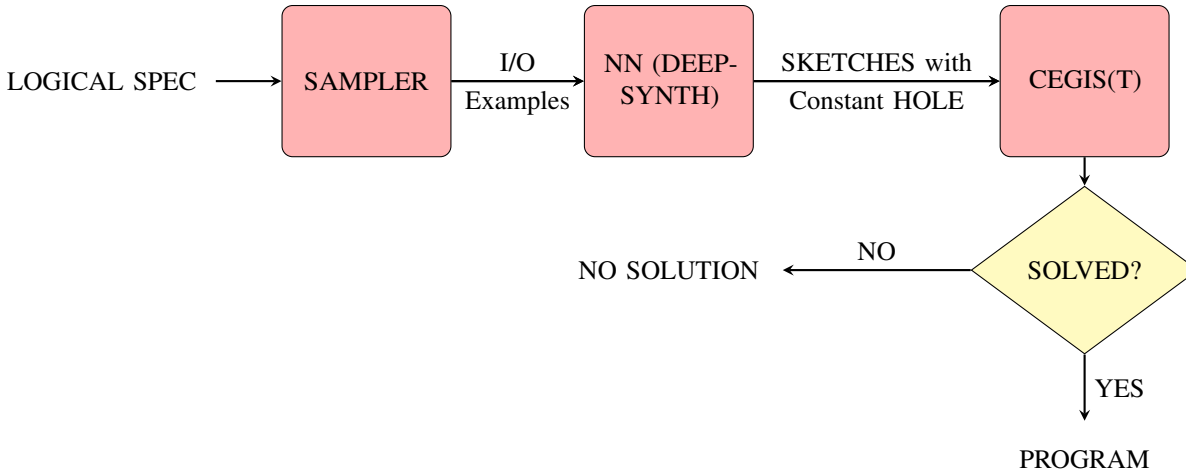


Fig. 6. SyGuS-Sketcher architecture

using logical embeddings and I/O embeddings and 3) A prototype tool SyGuS-Sketcher using existing techniques with experimentation plan.

REFERENCES

- [1] A. Abate, C. David, P. Kesseli, D. Kroening, and E. Polgreen. Counterexample guided inductive synthesis modulo theories. In *International Conference on Computer Aided Verification*, pages 270–288. Springer, 2018.
- [2] R. Alur, R. Bodik, G. Juniwal, M. M. K. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *2013 Formal Methods in Computer-Aided Design*, pages 1–8, 2013.
- [3] M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, and D. Tarlow. Deepcoder: Learning to write programs. *arXiv preprint arXiv:1611.01989*, 2016.
- [4] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 5(2):157–166, 1994.
- [5] Q. Chen, X. Wang, X. Ye, G. Durrett, and I. Dillig. Multi-modal synthesis of regular expressions. *arXiv preprint arXiv:1908.03316*, 2019.
- [6] Y. Chen, C. Wang, O. Bastani, I. Dillig, and Y. Feng. Program synthesis using deduction-guided reinforcement learning.
- [7] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [8] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [9] J. Devlin, J. Uesato, S. Bhupatiraju, R. Singh, A. rahman Mohamed, and P. Kohli. Robustfill: Neural program learning under noisy i/o, 2017.
- [10] Y. Feng, R. Martins, O. Bastani, and I. Dillig. Program synthesis using conflict-driven learning. *ACM SIGPLAN Notices*, 53(4):420–435, 2018.
- [11] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber, et al. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001.
- [12] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [13] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*, 2015.
- [14] M. Nye, L. Hewitt, J. Tenenbaum, and A. Solar-Lezama. Learning to infer program sketches. *arXiv preprint arXiv:1902.06349*, 2019.
- [15] E. Polgreen, R. Abboud, and D. Kroening. Counterexample guided neural synthesis, 2020.
- [16] X. Si, H. Dai, M. Raghothaman, M. Naik, and L. Song. Learning loop invariants for program verification. In *Advances in Neural Information Processing Systems*, pages 7751–7762, 2018.
- [17] A. Solar-Lezama. *Program Synthesis by Sketching*. PhD thesis, USA, 2008.
- [18] K. Wang and Z. Su. Learning blended, precise semantic program embeddings. *ArXiv*, vol. *abs/1907.02136*, 2019.

Learning Program Sketches from Logical Specifications using Reinforcement Learning

Ravi Raja
CSA, IISc
raviraja@iisc.ac.in

Abstract

This work aims at performing a literature review around Program Synthesis using Machine Learning techniques and exclusively writing reviews for the tools SKETCHADAPT and META-SOLVER. It also formulates a problem for synthesizing sketches for the General Track of SyGuS benchmarks and proposes a novel model META-SKETCHER aiming to solve it.

1 Introduction

Program Synthesis is an age-old problem of generating programs automatically that satisfies the user intent called Specification. It basically consists of three components: 1) Specification, 2) DSL, and 3) Search techniques. Many recent advances in this field are due to the fusion of Machine Learning tactics with it. Here, a brief overview of some interesting works are presented, that gives a basic understanding of the current state-of-the-art of the related area.

1.1 Literature Review

Deductive Synthesis : Some initial works on automatic program generation were based on Automated Theorem Provers [5, 8, 13]. The key idea followed by them was to construct the proof of the logical specification given by the user and then extract the logical program using it. Such approaches required the user to specify complete logical specifications which are often as complex as writing the programs itself. Also, it proved to be intractable due to huge program space. Syntax Guided Synthesis (SyGuS) [1] restricts the program space by defining a Domain Specific Language (DSL) with the help of a Grammar. 2 states the formal SyGuS problem statement.

Inductive Program Synthesis (IPS) : Inductive Program Synthesis (IPS) is based on inductive specifications (for e.g. I/O Examples, Demonstrations, Natural Language, or Partial Programs) emerged as an alternative. FlashMeta [10] by Polozov and Gulwani transforms the relationships between I/O pairs to intended programs. Along with syntactic specification it also includes the semantics of the language and follows a Divide and Conquer strategy to get the final program in real time. FlashFill [6] is a successful application of FlashMeta architecture used by Microsoft in spreadsheets. Program Synthesis by Sketching [12] by Armando Solar-Lezama synthesizes programs from partially complete programs called SKETCH.

Program Synthesis using Machine Learning : Menon et al. [7] uses the I/O examples to learn the probability distributions over productions in the Grammar converting it to Probabilistic CFG. This gives an order of the magnitude better performance over an enumerative baseline. The PCFG approach fails to include the contextual information stored in the partial programs generated at each instant while searching in the program space.

Neural Program Synthesis (NPS) : RobustFill by Devlin et al. [3] and DeepCoder by Balog et al. [2] are some recent advance in NPS. RobustFill designs a Neural Network that replaces the entire enumerative search part which is then trained to output a complete program. This may not guarantee the correctness of the final programs. Whereas DeepCoder uses a Neural Network to prune the search space rather than replacing the search part entirely.

META-SKETCHER : We propose META-SKETCHER, a novel approach for synthesizing programs from **logical spec.** with the help of **intermediate sketches**. As logical specs. can cover all possible program behavior and I/O specs. fails guarantee completeness. For ex. Given $f(x) = 2*x$, one can understand that the program output is twice the input, it cannot be derived correctly from I/O pair (1,2). It may get confused between addition of 1 and multiplication by 2. Therefore, this work primarily aims at solving SyGuS problems that appear in SyGuS Competition. Following two papers forms the basis of the work.

1.2 Review: Learning to Infer Program Sketches (SKETCHADAPT)

The key idea here is to generate program sketches from user spec. (I/O, natural language) that flexibly manages the work load between Neural Synthesis and Symbolic Search. **Sketches** are valid program tree in the DSL, where any number of sub-trees has been replaced by a special token called <HOLE>. Intuitively, this token designates locations in the program tree for which pattern-based recognition is difficult, and more explicit search methods are necessary. SKETCHADAPT [9] can be divided into two components: 1) **Neural Sketch Generator** which is a Seq2Seq neural architecture. It takes the spec. as input sequence and produces a good sketch as a sequence of tokens, and 2) **Program Synthesizer** which takes the generated sketch and outputs a final program by choosing grammar productions based on learned production probabilities. Generating a good sketch prunes out a large part of program space, thereby improving the synthesis time. Due to use of a symbolic search to fill up the sketches, programs produced are more generic. While this work is a huge success for the domains of List Processing and String Processing, it does not work at all on the kinds of problems found in the SyGuS-Comp where user intent is in the form of Satisfiability Modulo Theory (SMT) formulae. Further, it requires millions of training data for training the network which is not readily available for domains as in SyGuS.

1.3 Review: Learning a Meta Solver for Syntax-Guided Program Synthesis

This paper tries to formulate the problem of Program Synthesis in a RL framework. It gives following arguments to back it. As both the logical (semantic) spec and the grammar (syntactic spec) have complex structures and can vary from task to task, it poses significant challenges for learning across different tasks. Moreover, supervision is often unavailable for domain-specific synthesis tasks. Therefore, it's natural to formulate the problem in a Reinforcement Learning (RL) framework where the Environment is the Program Space specified by Grammar, Meta-Solver is the agent that performs an action by selecting a grammar production based on the reward returned by a SAT/SMT solver. The Meta-Solver [11] consists of three components: 1) an encoder, which embeds both the logical specification and grammar at the same time using a GGNN¹; 2) a grammar adaptive policy network which enables learning a transferable policy; and 3) a RL algorithm that jointly trains the embedding and adaptive policy with sparse reward². Comparing an Out-Of-Box setting which is trained and tested over same data with the two state-of-the-art classical synthesis tools CVC4 and EUSOLVER, it solves 129 and 153 out of 214 circuit synthesis problems respectively. While the Meta-Solver achieves speedup ranging from 2x up to 100x. Even though it performs really well in comparison to other state-of-the-art solvers such as CVC4 and EUSOLVER, the problem domain is too restrictive as it only solves the problems in Cryptographic circuit synthesis SyGuS benchmarks [4]. Also it does not consider generating sketches and combining a symbolic approach to complete it instead of fully depending upon RL.

¹Gated Graph Neural Network

²Rewards are called sparse when it is available only when the goal is reached (in this case a program is found)



Figure 1: META-SKETCHER Proposed Model

2 Problem Formulation

The Syntax-Guided Synthesis (SyGuS) problem is to synthesize a function f that satisfies two kinds of constraints:

- a syntactic constraint specified by a context-free grammar (CFG) G , and
- a semantic constraint specified by an SMT formula ϕ built from symbols in a background theory T along with f

This work aims at extending the work of [11] and generalizing it for SyGuS problems other than Cryptographic circuit synthesis problems by using the idea of sketches from [9].

Given a Dataset of N tasks $D = \{(\phi_i, G_i)\}_{i=1}^N$, the META-SKETCHER could have following components:

- Using Gated Graph Neural Network for learning the representation for the specification ϕ and the grammar G so that the synthesis process can utilize the structural information present in logical specification and grammar
- Learning a Policy $\pi_\theta : (\phi, G) \rightarrow \sigma$ where θ is the parameter of the learned policy and σ is the sketch or the partial program.
- Learning production probabilities p for the grammar productions i.e. for each non terminal α , probabilities should be assigned to each of its expansion $\alpha \rightarrow \beta_i$
- A Solver for solving sketches to get the final program using Symbolic Searches such as enumerative search guided by the production probabilities.

Two prime difficulties faced while solving such a problem are : 1) Scarcity of dataset: the dataset present in the SyGuS benchmarks are not enough to train a Neural Network, it requires millions of solved benchmarks, and 2) Sketches are not complete programs, so a SAT/SMT solver cannot verify it and hence cannot return a reward to the policy learner. These issues are addressed in following section.

3 META-SKETCHER

This section presents the META-SKETCHER model (see figure 1) that attempts to solve the three problems stated in section 2

3.1 Formal Definitions

Following are the formal definitions of some of the key terms.

Semantic Spec ϕ : These are formulas that consists of constraints written in the form of SMT formulae. It follows the same grammar G which is specified for the output program.

CFG G : A context free grammar (CFG) G is as defined in the [11] except the following modification.

The grammar should be modified for sketch generation by adding a new token $\langle \text{HOLE} \rangle$ which is used to replace any sub-tree in the program tree leaving that part to be searched using a symbolic search.

Output Sketch σ : The Output Sketch σ is a partial program generated using the grammar G . A good sketch will be the sketch that leads to final program f which satisfies both syntactic and semantic specifications.

Output Program f : The Output Program f is the final program that is generated from intermediate sketches. The final program f must satisfy both syntactic (grammar G) and semantic (SMT formula ϕ) specification

3.2 Novel Idea :

Synthesizing sketches for General SyGuS benchmarks where the specifications are in the form of logical formulae instead of input output examples using a Reinforcement Learning framework. The policy network should learn to output partial programs based on the rewards.

Rewards : As sketches are incomplete programs it should not be verified by SAT/SMT solvers with the given semantic specification, one needs to consider some indirect ways of receiving rewards. Using the learned production probabilities, likelihood of the sketches could be calculated, and this can easily serve the purpose.

There are no advancements regarding synthesizing sketches from SMT specifications and certainly there's no end-to-end synthesizing tool for SyGuS's General Track benchmarks ³ based on Reinforcement Learning setting. Therefore it would be interesting to see how well it works in comparison to existing tools.

4 Experiments

Experimental Setup : Following experiments are done on Intel(R) Xeon(R) CPU E5-2420 v2 @ 2.20GHz with 32 GB Memory

4.1 META-SOLVER

Experiments were performed to test the META-SOLVER on 30 out of 214 SyGuS Cryptographic circuit synthesis benchmark problems. It elegantly solved 22/30 problems in an average time of ~ 0.32 secs. While a state-of-the-art cvc4 ran out of memory for all the 30 problems. This indicates that the performance hike claimed by them are indeed correct.

It required more than 6 hours to conclude that a problem cannot be solved.⁴

The excel sheet for the results can be found at Excel Sheet for the Experiments performed

4.2 Preliminary Results

Preliminary experiments were performed in order to generate sketches using the existing framework of META-SOLVER by introducing <HOLE> in their grammar. Results show that it cannot handle recursive grammar and gives Recursion Error. A probable solution to this could be to use Monte Carlo Methods or Temporal-Difference Learning approaches. But it's testing remains to be performed in future.

5 Conclusion

This work does a decent Literature Review for the field of Program Synthesis and applicability of Machine Learning in it. It exclusively reviews the two recent tools for synthesizing programs, one with sketches and other without sketches namely SKETCHADAPT and META-SOLVER respectively. The study shows that no work has been done around generating sketches for SyGuS problems where the specification is in the form of SMT formulae. Hence, the META-SKETCHER model presented here could be looked upon for better synthesis results.

³<https://sygus.org/comp/2019/>

⁴The SyGuS 2017 competition gives each solver 4-core 2.4GHz Intel processors with 128 GB memory and wallclock time limit of 1 hour

6 Acknowledgements

I thank Stanly Samuel for his generous help in understanding the concepts from Programming Language community.

References

- [1] R. Alur, R. Bodik, G. Juniwal, M. M. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. *Syntax-guided synthesis*. IEEE, 2013.
- [2] M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, and D. Tarlow. Deepcoder: Learning to write programs. *arXiv preprint arXiv:1611.01989*, 2016.
- [3] J. Devlin, J. Uesato, S. Bhupatiraju, R. Singh, A.-r. Mohamed, and P. 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.
- [4] H. Eldib, M. Wu, and C. Wang. Synthesis of fault-attack countermeasures for cryptographic circuits. In *International Conference on Computer Aided Verification*, pages 343–363. Springer, 2016.
- [5] C. Green. Application of theorem proving to problem solving. In *Readings in Artificial Intelligence*, pages 202–222. Elsevier, 1981.
- [6] S. Gulwani. Automating string processing in spreadsheets using input-output examples. *ACM Sigplan Notices*, 46(1):317–330, 2011.
- [7] S. Gulwani, O. Polozov, R. Singh, et al. Program synthesis. *Foundations and Trends® in Programming Languages*, 4(1-2):1–119, 2017.
- [8] Z. Manna and R. J. Waldinger. Toward automatic program synthesis. *Communications of the ACM*, 14(3):151–165, 1971.
- [9] M. Nye, L. Hewitt, J. Tenenbaum, and A. Solar-Lezama. Learning to infer program sketches. *arXiv preprint arXiv:1902.06349*, 2019.
- [10] O. Polozov and S. Gulwani. Flashmeta: a framework for inductive program synthesis. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, pages 107–126, 2015.
- [11] X. Si, Y. Yang, H. Dai, M. Naik, and L. Song. Learning a meta-solver for syntax-guided program synthesis. 2018.
- [12] A. Solar-Lezama. Program synthesis by sketching. proquest. *Ann Arbor*, 2008.
- [13] R. J. Waldinger and R. C. Lee. Prow: A step toward automatic program writing. In *Proceedings of the 1st international joint conference on Artificial intelligence*, pages 241–252, 1969.