

Reinforcement Learning and Data-Generation for Syntax-Guided Synthesis

Julian Parsert^{1,2}, Elizabeth Polgreen²,

¹University of Oxford

²University of Edinburgh

julian.parsert@gmail.com, elizabeth.polgreen@ed.ac.uk

Abstract

Program synthesis is the task of automatically generating code based on a specification. In Syntax-Guided Synthesis (SyGuS) this specification is a combination of a syntactic template and a logical formula, and the result is guaranteed to satisfy both. We present a reinforcement-learning guided algorithm for SyGuS which uses Monte-Carlo Tree Search (MCTS) to search the space of candidate solutions. Our algorithm learns policy and value functions which, combined with the upper confidence bound for trees, allow it to balance exploration and exploitation. A common challenge in applying machine learning approaches to syntax-guided synthesis is the scarcity of training data. To address this, we present a method for automatically generating training data for SyGuS based on *anti-unification* of existing first-order satisfiability problems, which we use to train our MCTS policy. We implement and evaluate this setup and demonstrate that learned policy and value improve the synthesis performance over a baseline by over 26 percentage points in the training and testing sets. Our tool outperforms state-of-the-art tool *cvc5* on the training set and performs comparably in terms of the total number of problems solved on the testing set (solving 23% of the benchmarks on which *cvc5* fails). We make our data set publicly available, to enable further application of machine learning methods to the SyGuS problem.

Introduction

Syntax-Guided Synthesis (Alur et al. 2013) allows a user to synthesize code that is guaranteed to satisfy a logical and a syntactic specification, as shown in Example 1. The logical specification is given as a quantified formula: in the example, we must find a program such that $\forall x, y. f(x, y) \geq x \wedge f(x, y) \geq y \wedge f(x, y) = x \vee x(f, y) = y$. The syntactic specification is given as a context-free grammar. Applications of SyGuS include planning (Chasins and Newcomb 2016), deobfuscation (Jha et al. 2010), invariant synthesis (Fedyukovich et al. 2019), verified lifting (Kamil et al. 2016), and generating rewrite rules in SMT solvers (Reynolds and Tinelli 2017).

Surprisingly, the dominant techniques for SyGuS are still based on enumerative synthesis (Reynolds et al. 2019; Alur, Radhakrishna, and Udupa 2017), where the synthesizer exhaustively enumerates a grammar, and checks resulting pro-

grams against the specification until a correct solution is found. Good search heuristics can lead to extreme differences in performance (Reynolds et al. 2019), so one would expect machine learning techniques to be prominent.

In fact, in programming-by-example (PBE), where the specification is in the form of input-output (I/O) examples, the dominant heuristics are now machine-learning based (Balog et al. 2017; Bavishi et al. 2019). In PBE, training data is easy to generate by taking any program and generating I/O examples for it. We hypothesize that the reason these machine-learning techniques have not yet had a significant impact in SyGuS with logical specifications is the lack of easily available training data.

We present a synthesis algorithm that can learn its own heuristics via a reinforcement learning (RL) loop with relatively small quantities of training data. We use a Monte-Carlo Tree Search (MCTS) based approach to search for solutions within the syntactic specification using policy/value functions. The policy and value functions are iteratively improved via an RL loop. To evaluate a candidate program, we use a Satisfiability Modulo Theories (SMT) solver. Implementing these techniques we were able to achieve an improvement of over 26 percentage points (34.7% baseline vs 60.9% in the best iteration) on the testing set, and similar results on the training set. Compared with the state-of-the-art tool *cvc5* (Barbosa et al. 2022) we can solve 182 problems that *cvc5* cannot solve and on average 25 more problems on the training set while solving 6 fewer problems on the testing set.

To obtain training data, we develop a novel method for automatically generating SyGuS problems, based on anti-unification. Generating training data for logical specifications is challenging. In PBE, the standard approach to generating data is to randomly generate solutions and their corresponding input-output specifications. For example, we might generate the solution $(ite(\geq x y) x y)$ for f , where ite is the if-then-else operator, and the specification $f(2, 1) = 2$, $f(2, 2) = 2$, $f(1, 1) = 1$. This approach cannot be easily applied to SyGuS with logical specifications because, if we randomly generate solutions, it is extremely hard to infer a meaningful logical specification that captures that solution without also revealing the solution in the specification. This means that any inferred specifications are unlikely to be representative of real SyGuS problems. We could also randomly generate logical specifications, but these are highly unlikely

Listing 1: SyGuS problem, expressing the semantic constraint: $\forall x, y. f(x, y) \geq x \wedge f(x, y) \geq y \wedge (f(x, y) = x \vee f(x, y) = y)$, and a syntactic constraint.

```

;; The background theory
(set-logic LIA)
;; function to be synthesized
(synth-fun f ((x Int) (y Int)) Int
  ;; syntactic constraint
  ((S Int) (B Bool)) ;; non-terminals
  ((S Int) (x y 0 1
    (+ S S)(ite B S S)))
  (B Bool ((and B B) (or B B) (not B)
    (= S S) (>= S S))))
;; universally quantified variables
(declare-var x Int)
(declare-var y Int)
;; semantic constraints on the function
(constraint (>= (f x y) x))
(constraint (>= (f x y) y))
(constraint (or (= x (f x y))(= y (f x y))))
(check-synth)

```

to be feasible (i.e., to have solutions). Another approach used in the literature for invariant synthesis problems is to mutate existing specifications (Si et al. 2020), but 1) we must be careful not to mutate the specification in a way that renders the specification infeasible and 2) the approach in the literature does not generate training data with any new solutions, only new specifications.

In contrast, the novel method we present for generating training data for SyGuS generates new, meaningful logical specifications beyond PBE. We use syntactic *unification* (where assignments to variables are found to make two expressions equal) and *anti-unification* (where a generalization of two terms is found) techniques on existing SMT problems to generate feasible and interesting logical specifications. This, as opposed to mutating existing specs randomly, generates feasible problems with new solutions.

The main contributions of this paper are:

- We frame SyGuS as a tree search problem so that we can use a Monte-Carlo tree search (MCTS) based synthesis algorithm that uses machine learned policy and value predictors trained in a reinforcement learning loop.
- We train policy and value predictors in an RL loop.
- We present a method for generating SyGuS problems from first-order SMT problems using anti-unification and unification to overcome a lack of training data.
- We evaluate our method on a combination of pre-existing benchmark sets from the SyGuS competition (Alur et al. 2017) and our newly generated problem set and compare it to the state-of-the-art SyGuS solver cvc5.

Related work: Machine-learning based synthesis methods have been applied to PBE program synthesis problems (Balog et al. 2017; Bavishi et al. 2019; Morton et al. 2020; Bunel et al. 2018; Odena et al. 2021) where the specification of correctness is a set of input-output examples, and strategy synthesis (Medeiros, Aleixo, and Lelis 2022). Two approaches in literature use RL. The first uses a pre-trained policy that

is then updated using deductive reasoning guided RL (Chen et al. 2020). The second applies RL training an agent to interact with a Read-Eval-Print loop (Ellis et al. 2019). Neither of these approaches are applied to synthesis with logical specifications, and so neither offers solutions to the challenges of a relatively sparse training data set.

Common algorithms for solving SyGuS are based on Oracle Guided Inductive Synthesis (OGIS) (Jha and Seshia 2017; Solar-Lezama 2009), which alternates between a search phase that enumerates the space of possible programs with various heuristics (Reynolds et al. 2019; Feng et al. 2018; Alur, Radhakrishna, and Udupa 2017; Abate et al. 2018; Lee et al. 2018), and an oracle that returns feedback on candidate programs. Previously, traditional MCTS has been applied without learned guidance to expression search (Cazenave 2013), and with learned guidance to program search (Toni, Erculiani, and Passerini 2021). The closest to our work is Euphony (Lee et al. 2018), which learns a probabilistic grammar and uses this to guide an A* style search algorithm. Euphony is pre-trained and requires a corpus of training data problems with solutions, whereas we learn through RL and do not require solutions in our training data.

RL (Laurent and Platzter 2022) and graph neural networks (Si et al. 2020) have been applied to invariant synthesis, which is a version of synthesis with logical specifications, without syntactic templates. This domain also suffers from training data scarcity and data-driven approaches have to consider ways to overcome this. For instance, (Si et al. 2020) take existing loop invariant specifications and mutate them in ways that are guaranteed to keep the solution the same.

The algorithm that we adapt to enumerative synthesis is based on AlphaZero (Silver et al. 2018) and has successfully been applied to first-order logic theorem proving (Kaliszyk et al. 2018) and combinator synthesis (Gauthier 2020).

Anti-unification was first presented by Plotkin (Plotkin 1970). It is the dual of unification, which is widely used in theorem proving, logic programming, and term re-writing (Kutsia 2013; Yernaux and Vanhoof 2019; Baader and Nipkow 1998). Anti-unification has been applied to solving PBE problems (Raza, Gulwani, and Milic-Frayling 2014).

Background

Syntax-guided synthesis (SyGuS): A SyGuS problem, see Example 1, is a tuple $\langle \tau, f, \phi, G \rangle$ where τ is a background theory, f is the function to be synthesized, ϕ is a quantifier-free formula and G is a context-free grammar. The task is to find a body for f such that f is in the language specified by G and the formula ϕ is τ -valid, i.e., the formula $\exists f \forall \vec{x}. \phi(f)$ must be true where \vec{x} is the vector of free variables in ϕ .

Context-free grammars: A grammar G is a tuple $G = (N, T, R, S_t)$, where N is a finite set of symbols (non-terminals); T is a finite set of symbols (characters of the language); N and T are disjoint; R is a set of production rules, where each rule is of the form $N \rightarrow (N \cup T)^*$, where $*$ represents the Kleene star operation; and S_t is a symbol in N (the start symbol). We use L^G to denote the language specified by the grammar G .




Figure 1: Grammar and corresponding grammar tree.

Given a context-free grammar $G = (N, T, R, S_t)$, the grammar tree is defined as follows: S_t is the only root (i.e. node with no parent), and for every node r , $r[\alpha \rightarrow \beta]$ is a child of r if $\alpha \in \text{vars}(r)$ and $\alpha \rightarrow \beta \in R$. Note that $r[\alpha \rightarrow \beta]$ is the substitution of β for α for the leftmost occurrence of a non-terminal variable α in r .

Figure 1 shows an example of a *grammar tree*. Every valid parse tree according to the grammar is a path from the root node to a leaf node of the grammar tree. Hence, every leaf node denotes a complete program and if the language of the grammar is infinite the resulting tree is also infinite.

Monte-Carlo Grammar Tree Search

We describe enumerative function synthesis as a single-player game. Let $P = \langle \tau, f, \phi, G \rangle$ be a SyGuS problem. Each game state consists of a static component, the specification ϕ which does not change throughout the game, and a dynamic component which corresponds to the vertices in the grammar tree of G . Hence, every state $s \in S$ of the synthesis game is a pair $\langle \phi, H \rangle$ where H is a partial program generated by G . The set of actions A that can be taken in a state solely depends on the dynamic component – the expression H . If H is a complete program we have a final state where no further actions can be taken and we use an SMT solver to check if H is a solution to P . If we have not reached a final state the possible actions correspond exactly to all leftmost reductions for H . Thus, the game tree corresponds to the grammar tree of the grammar G where the root is the state $\langle \phi, S_t \rangle$ with S_t being the start symbol of G . For all other nodes, we have that if $s = \langle \phi, H \rangle$ is a node in the game then $s' = \langle \phi, H' \rangle$ where H' can be obtained from H by one application of a grammar rule in G is a child node of s .

We now present the search algorithm as an agent playing the synthesis game to find solutions. In subsequent sections we discuss how we incorporate machine learning with a balancing of exploration and exploitation to guide this agent.

Algorithm 1: Big-Steps

```

Data: SyGuS problem  $P = \langle \tau, f, \phi, G \rangle$ 
Result: fail or solution to  $P$ 
act_nd  $\leftarrow$  root( $GT(G)$ );
for  $i \leftarrow 0$  to MAX.BIGSTEPS do
     $r \leftarrow$  rollout(active_nd,  $P$ );
    if  $r$  is solution then
        return  $r$ ;
    else
        act_nd  $\leftarrow$ 
            most_visited_child(act_nd,  $GT(G)$ )
    end
end
return fail

```

Algorithm

Similar to Kaliszky et al (Kaliszky et al. 2018), we adapt the Monte-Carlo based tree search used in AlphaZero (Silver et al. 2018). The algorithm consists of four main phases: *Big-Steps*, *Rollout*, *Expansion*, and *Backpropagation*. The algorithm’s input is a SyGuS problem $P = \langle \tau, f, \phi, G \rangle$ and the output is either a solution $f \in L^G$ that satisfies the constraints or a general failure (i.e. timeout). During the search we keep a record of the visit count $N(s)$ with the default visit count 0 as well as other statistics that we will discuss later. During backpropagation, these records are updated. We also keep track of whether or not a node has been expanded. In detail, the four phases work as follows:

Big-Steps: As the outermost loop of the search procedure, we start at the root of the grammar tree by setting the *active node* (act_nd in the algorithm) to the root of $GT(G)$ as depicted in Algorithm 1. Starting from the active node we execute the other phases of the algorithm which among other things update the visit counts of all nodes. If the rollout finds a solution, we are finished and return the solution. Otherwise, we update the active node to the most visited immediate child for the next iteration, i.e., the child node with the highest visit count.

Rollout, Expansion, and Backpropagation: In this phase we perform rollouts from a given starting node in the grammar tree, i.e., traversing the grammar tree starting from that node. The purpose of this process is to collect extensive visit count data by applying *expansion* steps, *checking potential solutions* where appropriate, and applying *backpropagation* so that Algorithm 1 can make the best decision possible in choosing the next active node. As shown in Algorithm 2 we follow a path in the grammar tree starting from a node called active_nd to the first unexpanded node by always selecting the best successor of the current node. Later, we derive a notion of “best” to improve the search guidance, using the upper confidence bound on the values of the nodes. If the first unexpanded node is a leaf node we take the complete function it represents and the *verify* function uses an SMT solver to check if the function is a solution to P . If it is, we return the solution, otherwise, we set the value (i.e. reward) of the node to 0. If the node is not a leaf node, we expand the

Algorithm 2: Rollout

```

Data: active_nd and  $P$ 
Result: fail or solution to  $P$ 
for  $i \leftarrow 0$  to  $\text{MAX\_ROLLOUTS}$  do
  current_nd  $\leftarrow$  active_nd;
  sub_path = [current_nd];
  while is_expanded(current_nd) do
    /* select best successor node based on UCT */
    current_nd  $\leftarrow$ 
      best_successor(current_nd);
    /* append the current node to the path */
    append(sub_path, current_nd)
  end
  if isLeaf(current_nd) then
    if verify( $P$ , current_nd) then
      return current_nd;
    else
       $W(\text{current\_nd}) \leftarrow 0$ ;
      break;
    end
  else
    expand(current_nd,  $P$ )
  end
  /* add the value of the current node to all nodes on the
  path and increment the visit counts */
  backpropagate(sub_path,  $W(\text{current\_nd})$ )
end

```

node by setting the visit count to 1. Finally, the *backpropagation* function increments the visit counts by 1 and adds the value of the current node to all nodes along the path we have taken. This whole process is repeated MAX_ROLLOUT times. In the experiments, we also account for a decay factor that decreases with every big step. Since we terminate the search upon successful verification (i.e. we found a correct program) backpropagating a positive reward is not necessary.

Guiding the Search

In this section, we introduce concepts for the *quality of nodes and actions* as well as a method to balance the exploitation of known – with the exploration of unknown subtrees.

Similar to AlphaZero, we use the notions of value and policy of each state and edge in the game tree of the synthesis game. The value function $\nu(s) : S \mapsto \mathbb{R}$ maps each state in the grammar tree $s \in S$ to a real number that represents the “quality” of the state s . This estimates the likelihood of reaching a successful state in the subtree starting from the given state. The second heuristic that we use is the policy function $\pi(s, a) : S \times A \mapsto \mathbb{R}$ mapping a state s in the grammar tree and action a to a real number, which estimates the probability of success when committing to action a from state s . For readability we use the notation $\pi(s, s')$ where $s, s' \in S$ to denote the policy of the action a that leads from state s to s' . In addition to the visit count as mentioned above we also keep a record of the cumulative value $W(s)$ and policy $\pi(s, s')$ for each state $s, s' \in S$ in the search graph these values are queried and saved during the *expansion*

phase. During the backpropagation phase, we also add the value $\nu(\text{end})$ of the last node end in the path to all other nodes resulting in the *cumulative* value.

Upper Confidence Bound for Trees: If we always select the subtrees with the highest average values and policies we would never explore unseen parts of the tree and gain more information about other parts of the graph. Moreover, if the value and policy are not perfect (i.e. bias against subtrees that contain the solution) we might end up repeatedly exploring a path that does not lead to a successful node. Conversely, if we only look for unseen paths in the search tree we do not use the information provided by the value and policy functions rendering them useless. Hence, our goal is to strike a balance between *exploration* of unknown parts of the graph and *exploitation* of the information we have already obtained in order to reach a solution as fast as possible. To this end, we use the upper confidence bound for trees (UCT) (Kocsis and Szepesvári 2006) as a heuristic to select the “best” child node, whilst balancing the trade-off between exploration and exploitation. Given a parent node $p \in S$ and child node $c \in S$ the UCT value of the child node when starting in the parent node is calculated as follows:

$$\text{uct}(p, c) = \frac{W(c)}{N(c)} + \gamma * \pi(p, c) * \sqrt{\frac{\log N(p)}{N(c)}}$$

This is the sum of two terms: an exploitation term $\frac{W(c)}{N(c)}$ which calculates the average value per visit of the child node and an exploration term $\pi(p, c) * \sqrt{\frac{\log N(p)}{N(c)}}$. The constant γ defines the relation between exploration and exploitation (i.e. if $\gamma = 0$, we only consider exploitation). The exploration term weighs the policy $\pi(p, c)$ (i.e. likelihood of success when choosing the action that leads from p to c) with a term that increases in value when the visit count for p increases and the visit count for c stagnates. Conversely, if the child c visit count increases comparatively to the parent visit count the term decreases in value, and thus the exploration value decreases. In combination, this leads to child nodes with fewer visits than their siblings having a larger exploration value.

Learning Policy and Value Functions

Features: To facilitate the use of machine learning models we have to design vector representations for synthesis states and actions. The main features we use are term walks which are parent/child pairs in the syntax tree of a term. These syntactic features are similar to those used in first-order logic theorem provers (Jakubuv and Urban 2017; Chvalovský et al. 2019; Färber, Kaliszzyk, and Urban 2021; Kaliszzyk et al. 2018). We use a bag-of-words representation where the words are the hashes of all constants, variables, and term walks of size two. This results in a feature vector of size n , where the entry at the i th position denotes the number of features with hash congruent to $i \bmod n$. The constant n is the hash base and describes the maximum number of distinct features.

Finally, we apply this method to the formula ϕ and term H to obtain two feature vectors the concatenation of which is the feature vector of the state $\langle \phi, H \rangle$. Similarly, we encode

the action a that leads from state $s = \langle \phi, H \rangle$ to $s' = \langle \phi, H' \rangle$ as the concatenation of the feature vectors of H and H' . Thus, the feature vector for policy estimation is the concatenation of the feature vectors for ϕ , H , and H' .

Learning: Policy and value prediction can be described as a regression model, so many machine learning models can be used here. However, due to the speed of tree traversal and node selection in our algorithm, the calls to the learning models are the bottlenecks, much like in theorem proving (Chvalovský et al. 2019; Irving et al. 2016). Hence, choosing models with low prediction latency is vital. In our experiments, we use gradient boosted trees provided by XGBoost (Chen and Guestrin 2016), which have been successful in theorem proving. We will mention more implementation details later.

Reinforcement Learning: We use an RL loop to train the policy and value models. In the first iteration, we run the synthesis procedure on all training problems using a default policy of 1 for each action and a default value of $0.95^{\#(NT)}$ where $\#(NT)$ is the number of non-terminals in the synthesized expression. Once completed, we collect the following training data from the problems and their solutions:

Policy training data: For each pair of consecutive states $s = \langle \phi, H \rangle$ and $s' = \langle \phi, H' \rangle$ in each successful search path we take the pair $(\langle \phi, H, H' \rangle, p)$ where p is the number of visits of s' in relation to the sum of all visits of all children of s .

Value training data: For each state $s = \langle \phi, H \rangle$ on every path we take the pair $(\langle \phi, H \rangle, 0.9^D)$ where D is the distance from the last (i.e. successful) node as training data. In case of a failed search, we use the pair $(\langle \phi, H \rangle, 0)$.

At the end of every RL iteration, this data together with the data of previous iterations is used to train new policy and value functions that will be used for the next iteration.

Generating SyGuS Problems

The algorithm above is trained using an RL loop that learns policy and value functions. Consequently, the larger the corpus of training data, the better the algorithm is likely to perform. Randomly generating SyGuS problems as data is not practical: if we randomly generate the specification ϕ , it is highly unlikely that a valid f exists; and if we randomly generate the solution f , we need to find a way to infer a meaningful specification ϕ that admits f and does not give away the answer. The number of benchmarks available in the linear arithmetic (LIA) category of the SyGuS competition is relatively small (less than 1000), due to the relative immaturity of the field, but there are many more first-order LIA SMT problems, so we now describe a novel technique for generating SyGuS problems from SMT problems, using syntactic unification (Martelli and Montanari 1982) and anti-unification (Cerna and Kutsia 2023; Plotkin 1970). The former finds variable assignments to equations that make the terms syntactically equal. For example when applying unification to the equation $(5 \odot x) \bullet (3 \ominus c) = (y \odot 1) \bullet z$ we get the unifier $\{x \mapsto 1, y \mapsto 5, z \mapsto (3 \ominus c)\}$. If we apply the mapping of variables to the term above we get the two terms that are syntactically equal $(5 \odot 1) \bullet (3 \ominus c) = (5 \odot 1) \bullet (3 \ominus c)$.

The latter is a process that produces a generalization common to a set of symbolic expressions. For example, applying anti-unification to the two terms $(2 \bullet 1) \oplus 8$ and $(1 \bullet 3) \oplus 5$ we obtain the least general generalizer (LGG) $x \bullet y \oplus w$. In other words, we obtain a term such that if we apply unification with either of the original terms we obtain a unifier that turns the LGG into that original term. Given a τ -valid SMT problem $Q = \langle \tau, \phi \rangle$, i.e., an SMT problem that is true for all variable assignments, we generate a SyGuS problem P as follows:

1. Heuristically select a set S of sub-terms of Q
2. Compute LGG l of S with fresh variables x_1, \dots, x_n
3. Replace each term $t \in S$ in Q with synthesis target $f(x_1, \dots, x_n)$ using unifier of $t = l$ to determine the correct variable mapping
4. Generate a grammar G that produces all terms in τ using variables x_1, \dots, x_n .
5. Create SyGuS problem $P = \langle \tau, f, \phi, G \rangle$.

The LGG is one possible solution to the resulting problem P . The first step has the largest impact on the quality of the resulting problem as, if the wrong subterms are chosen, the resulting LGG might be trivial (e.g. for the terms $\{5+4, 9*x\}$ the LGG is basically the identity function). In our case, we search for sub-terms that have the largest LGG as well as sub-terms of type `int`. Sub-terms of type `bool` often lead to trivial solutions such as `True` or `False`.

Analysis of New Training Data

We use SMT problems taken from the LIA and QFLIA (Quantifier-free Linear Integer Arithmetic) tracks of the SMT competition (Weber et al. 2019) to generate new SyGuS problems. We run `cvc5` (Barbosa et al. 2022) on the new problems to group them into the following categories, in approximate order of complexity: **Basic (B) problems** that `cvc5` can solve and where a valid solution is a function returning a single constant or variable); **Straight-line (S) problems** that `cvc5` can solve and are not basic but where a valid solution does not need control flow (i.e., there is a valid solution that performs simple mathematical operations); **control-flow (C) problems** that `cvc5` can solve and where any valid solution needs control flow (i.e., a valid solution must contain if then else statements); and **unsolved (U) problems** that `cvc5` cannot solve with a 120s time-out.

Table 1 reports the number of problems in each of these categories in our *new* data-set and in the original (*old*) data-set from the SyGuS competition. The SyGuS competition data comprises all benchmarks in LIA from the General and Invariant Synthesis tracks with a single synthesis function. If a benchmark has no grammar, we augment the benchmark with the same grammar as our generated data.

After running our data generation procedure overnight, we generate 8186 new SyGuS LIA problems. We remove basic problems and filter the remaining problems to reduce the number of very similar problems, using a heuristic based on the file names and duplicate comments as these are usually indicative of the source of the SMT benchmark, and correlate with very similar problems. This gives 944 new SyGuS LIA problems, and, in total, we have a problem set consisting of 1901 SyGuS LIA problems. This more than doubles the total

Data set	#B	#S	#C	#U	Total
SyGuS-comp (old)	290	93	38	536	957
Generated (new)	3760	2495	1693	220	8168
Filtered new (new)	–	211	513	220	944

Table 1: Number of problems by category: basic problems (B), straight-line problems (S), problems with control-flow (C), and unsolved problems (U).

number of SyGuS LIA problems, and increases the number of solvable benchmarks with control flow by $10\times$. The number of problems we generate is primarily limited by the time that we run the data generation, as multiple problems can be generated from each first-order problem, and the number of publicly available SMT problems increases year-on-year year.

Experimental Evaluation

We implemented the Monte-Carlo tree search algorithm with the learned policy and value using RL as discussed in previous sections. Our experimental evaluation seeks to answer the following questions:

Q1: Can MCTS with RL be used in function synthesis?

Q2: Does our data generation improve the performance?

Q3: How does MCTS compare to other techniques?

The implementation is written in C++ using Z3 (de Moura and Bjørner 2008) as an SMT solver to check the correctness of potential candidate solutions. We use XGBoost (Chen and Guestrin 2016) as a library to provide gradient-boosted trees for policy and value estimation. Running the experiments multiple times with different hyper-parameters ranging from (10 to 30) we found that the best tree depths are 20 and 25 for value and policy. We did not experience any significant improvements by changing other hyper-parameters. In the MCTS we do 30 big-steps and 6500 rollouts with a decay factor of 0.98^B where B is the number of big-steps. These were optimized for a 100s timeout. As a hash base for the feature vectors, we use $2^{12} - 3 = 4093$. Our synthesis algorithm exclusively learns from its previous iterations without external guidance or solutions. For training, we take data obtained from the previous 4 iterations instead of all previous iterations (we tried values between 2 and 20). We do not train during runs on the test set. Previously, we had a Python-based implementation where we experimented with other machine learning models such as K-Nearest Neighbours and Linear Regression as well as larger hash bases, and found that tree models outperformed other machine learning models.

Benchmarks & Setup: We use the SyGuS LIA data set consisting of 1901 problems described in the previous section. We create a 75 : 25 training/testing split on the data sets for each of the experiments. We make the data set as well as the code with Dockerfile and instructions on how to install dependencies, compile the code, and run experiments available¹. The experiments were run on Amazon Elastic Compute

Cloud (EC2) on a `r6a.8xlarge` instance with an AMD EPYC 7R13 CPU and 256GB of Memory running Amazon Linux 2. We run each search procedure with a timeout of 100 seconds. Our implementation introduces non-determinism in two ways: training of the machine learning models and a random tie break in child selection when two or more child nodes have the same scores. We ran experiments 5 times with fixed seeds.

Results

Q1: Can MCTS with RL be used in function synthesis?

We run experiments on the union of all benchmark sets resulting in 1425 training and 476 testing problems. The results are shown in experiment 1 (Exp. 1) of Table 2. The baseline is the first iteration of the reinforcement learning loop. In this iteration, we have not gathered any training data and have no trained guidance so use default value and policy functions. From the second iteration onwards, policy and value functions are trained on the data obtained from the training set in the previous iterations (no data is collected from the testing set). Table 2 shows a comparison of the performance of the best iterations of each of the 5 experiment runs on the training and testing set. The best iteration of each experiment on average solves 859.8 and 289 problems on the training and testing set respectively. This is an improvement of 27.2 and 26.4 percentage points compared with the first iteration for training and testing sets. Figure 2 plots the number of solved testing problems for each iteration for each experiment.

Q2: Does our data generation improve performance?

Doing a more fine-grained analysis of the experiments, we find that the testing set in the previous experiments contains 249 problems from the old data set. Of these, between 75 (30.1%) and 78 (31.3%) are solved with a mean of 76.4 (30.7%) and a stdev of 1.14. To compare, we conduct two additional experiments with the following data set-ups, with data sets from *Old* and *New* problems, as previously described, and report the results in Table 2

1. Train with *old* and test with *old* problems (Exp. 2)
2. Train with *new* and test with *old* problems (Exp. 3).

In Exp. 2 the trained search agent solves almost 1.85 times as many problems as the baseline (16.6% baseline and 30.7% on average in the best iteration on the testing set). In contrast, Exp. 3 shows that training exclusively on new problems and testing exclusively on old problems only improves performance on the testing set (i.e. old problems) by 2.4% (15.9% on baseline and 18.3% on average in the best iteration). In line 8 we highlight the performance of the setup in Experiment 3 on the same test set as in Experiment 2, which has a similar marginal performance improvement. We find that the new training data does not improve the performance on the old test set compared with simply training on the old data, but as can be seen in Exp. 1 (line 3) of Table 2, it also does not diminish the performance on the old problems, whilst still enabling us to solve a significant number of the new problems that `cvc5` is unable to solve. We hypothesize this is because our new benchmark set is not typical of the existing SyGuS benchmarks (this is also demonstrated by `cvc5`'s failure to solve many of these examples).

¹<https://zenodo.org/records/10377451>

	action	data	total	baseline		Solved in Best Iteration					cvc5		
				#	#	%	min	max	# mean	% mean	stdev	#	%
Exp. 1	1	train	old + new	1425	471.6	33.1%	852	865	859.8	60.3%	3.97	834	58.5%
	2	test on	old + new	476	163.4	34.3%	287	292	289	60.7%	1.87	295	62.0%
	3	test on	old	249	49.4	19.8%	75	78	76.4	30.7%	1.14	115	46.2%
Exp. 2	4	train on	old	717	112.2	15.7%	205	224	210.8	29.4%	7.6	315	41.0%
	5	test on	old*	240	39.8	16.6%	69	79	73.6	30.7%	3.58	108	45.0%
Exp. 3	6	train on	new	944	479.6	50.8%	874	886	879.6	93.2%	3.37	727	77.0%
	7	test on	old	957	151.8	15.9%	169	184	175.4	18.3%	5.97	402	42.0%
	8	test on	old*	240	38.6	16.1%	45	50	47.4	19.8%	2.51	108	45.0%

Table 2: Results summary. Training and testing sets are disjoint sets. The two old* data sets are the exact same sets.

Q3: How does MCTS compare to other techniques? We compare our approach to cvc5 (Barbosa et al. 2022), the SyGuS solver which performed best in the most recent SyGuS competition (Alur et al. 2017). In Exp. 1 the best iteration of our tool on average solves 25.8 problems more (min. 18, max. 31 more) than cvc5 in the training set and on average 6 fewer (min. 3, max. 8 fewer) on the test set. On average, we solve 41 test set and 141 training set problems on which cvc5 fails. Exp. 2 and 3 show that cvc5 performs poorly compared with our synthesis tool on the newly generated benchmarks. In particular, when used as training set, we solve 93.2% of the new problems while cvc5 only solves 77%. In contrast, the best iteration of our tool lags behind cvc5 by around 12 to 15 percentage points on the old data set, as shown by Exp. 2. cvc5 is a mature software consisting of over 400k lines of code incorporating many decision procedures and domain-specific heuristics to improve function search. In comparison, our implementation exclusively uses the learned search heuristics presented as described in previous sections, and is able to perform comparably to cvc5 on the testing set and even outperform cvc5 in some setups.

Limitations We do not compare our approach to specialized invariant synthesis tools, which may outperform our technique on invariant benchmarks but cannot handle the general SyGuS problems. We are also unable to compare to Euphony (Lee et al. 2018) as the available code does not compile, and they do not report results on LIA. Our conclusions are based on results on the SyGuS datasets, using hyperparameters tuned on these datasets, and these results may not hold in other domains. We also require a base set of problems to be solved in the first iteration to generate data to train from in succeeding iterations. Finally, our data-generation approach is only applicable to domains where first-order (verification) problems are available, and we wish to solve second-order (synthesis) problems. However, this is the case in many areas such as software verification/synthesis (David et al. 2018), unit test generation, and reactive motion planning.

Conclusions and Future Work

We presented a synthesis algorithm for SyGuS based on MCTS, incorporating machine-learned policy and value pre-




Figure 2: Number of solved problems in the testing set for each iteration in each of the 5 experiment runs. Each line represents one experimental run.

dictors and guidance based on UCT, and developed a method to generate SyGuS problems from preexisting SMT problems using anti-unification and unification. Our approach, using simple syntactic features and gradient-boosted tree models, improves on the baseline synthesizer by over 26% on training and testing sets in the best iteration, and solves benchmarks that are out of reach of state-of-the-art synthesis tools. In future work, we will exploit our new data to investigate the application of other data-dependent machine learning methods to the SyGuS domain, for instance, algorithm selection and run-time prediction (Healy, Monahan, and Power 2016), and we hope that others will do the same.

Acknowledgments

This work was in part supported by the Oxford-DeepMind Graduate Scholarship, the Engineering and Physical Sciences Research Council (EPSRC), an Amazon Research Award, and the Austrian Science Fund (FWF) project AUTOSARD (36623).

References

- Abate, A.; David, C.; Kesseli, P.; Kroening, D.; and Polgreen, E. 2018. Counterexample guided inductive synthesis modulo theories. In *International Conference on Computer Aided Verification*, 270–288. Springer.
- Alur, R.; Bodík, R.; Juniwal, G.; Martin, M. M. K.; Raghothaman, M.; Seshia, S. A.; Singh, R.; Solar-Lezama, A.; Torlak, E.; and Udupa, A. 2013. Syntax-guided synthesis. In *Formal Methods in Computer-Aided Design, FMCAD 2013, Portland, OR, USA, October 20-23, 2013*, 1–8. IEEE.
- Alur, R.; Fisman, D.; Singh, R.; and Udupa, A. 2017. Syntax Guided Synthesis Competition. <https://sygus-org.github.io>. Accessed: 2024-01-16.
- Alur, R.; Radhakrishna, A.; and Udupa, A. 2017. Scaling Enumerative Program Synthesis via Divide and Conquer. In *TACAS (1)*, volume 10205 of *Lecture Notes in Computer Science*, 319–336.
- Baader, F.; and Nipkow, T. 1998. *Term rewriting and all that*. Cambridge University Press.
- Balog, M.; Gaunt, A. L.; Brockschmidt, M.; Nowozin, S.; and Tarlow, D. 2017. DeepCoder: Learning to Write Programs. In *ICLR (Poster)*. OpenReview.net.
- Barbosa, H.; Barrett, C. W.; Brain, M.; Kremer, G.; Lachnitt, H.; Mann, M.; Mohamed, A.; Mohamed, M.; Niemetz, A.; Nötzli, A.; Ozdemir, A.; Preiner, M.; Reynolds, A.; Sheng, Y.; Tinelli, C.; and Zohar, Y. 2022. cvc5: A Versatile and Industrial-Strength SMT Solver. In *TACAS (1)*, volume 13243 of *Lecture Notes in Computer Science*, 415–442. Springer.
- Bavishi, R.; Lemieux, C.; Fox, R.; Sen, K.; and Stoica, I. 2019. AutoPandas: neural-backed generators for program synthesis. *Proc. ACM Program. Lang.*, 3(OOPSLA): 168:1–168:27.
- Bunel, R.; Hausknecht, M.; Devlin, J.; Singh, R.; and Kohli, P. 2018. Leveraging grammar and reinforcement learning for neural program synthesis. In *International Conference on Learning Representations*.
- Cazenave, T. 2013. Monte-Carlo Expression Discovery. *Int. J. Artif. Intell. Tools*, 22(1).
- Cerna, D. M.; and Kutsia, T. 2023. Anti-unification and Generalization: A Survey. *CoRR*, abs/2302.00277.
- Chasins, S. E.; and Newcomb, J. L. 2016. Using SyGuS to Synthesize Reactive Motion Plans. In *SYNT@CAV*, volume 229 of *EPTCS*, 3–20.
- Chen, T.; and Guestrin, C. 2016. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, 785–794. New York, NY, USA: ACM. ISBN 978-1-4503-4232-2.
- Chen, Y.; Wang, C.; Bastani, O.; Dillig, I.; and Feng, Y. 2020. Program Synthesis Using Deduction-Guided Reinforcement Learning. In *CAV (2)*, volume 12225 of *Lecture Notes in Computer Science*, 587–610. Springer.
- Chvalovský, K.; Jakubuv, J.; Suda, M.; and Urban, J. 2019. ENIGMA-NG: Efficient Neural and Gradient-Boosted Inference Guidance for E. In *CADE*, volume 11716 of *Lecture Notes in Computer Science*, 197–215. Springer.
- David, C.; Kesseli, P.; Kroening, D.; and Lewis, M. 2018. Program Synthesis for Program Analysis. *ACM Trans. Program. Lang. Syst.*, 40(2): 5:1–5:45.
- de Moura, L. M.; and Bjørner, N. S. 2008. Z3: An Efficient SMT Solver. In *TACAS*, volume 4963 of *Lecture Notes in Computer Science*, 337–340. Springer.
- Ellis, K.; Nye, M. I.; Pu, Y.; Sosa, F.; Tenenbaum, J.; and Solar-Lezama, A. 2019. Write, Execute, Assess: Program Synthesis with a REPL. In *NeurIPS*, 9165–9174.
- Färber, M.; Kaliszzyk, C.; and Urban, J. 2021. Machine Learning Guidance for Connection Tableaux. *J. Autom. Reason.*, 65(2): 287–320.
- Fedyukovich, G.; Prabhu, S.; Madhukar, K.; and Gupta, A. 2019. Quantified Invariants via Syntax-Guided Synthesis. In *CAV (1)*, volume 11561 of *Lecture Notes in Computer Science*, 259–277. Springer.
- Feng, Y.; Martins, R.; Bastani, O.; and Dillig, I. 2018. Program synthesis using conflict-driven learning. In *PLDI*, 420–435. ACM.
- Gauthier, T. 2020. Deep Reinforcement Learning for Synthesizing Functions in Higher-Order Logic. In *LPAR*, volume 73 of *EPiC Series in Computing*, 230–248. EasyChair.
- Healy, A.; Monahan, R.; and Power, J. F. 2016. Predicting SMT Solver Performance for Software Verification. In *F-IDE@FM*, volume 240 of *EPTCS*, 20–37.
- Irving, G.; Szegedy, C.; Alemi, A. A.; Eén, N.; Chollet, F.; and Urban, J. 2016. DeepMath - Deep Sequence Models for Premise Selection. In *NIPS*, 2235–2243.
- Jakubuv, J.; and Urban, J. 2017. ENIGMA: Efficient Learning-Based Inference Guiding Machine. In *CICM*, volume 10383 of *Lecture Notes in Computer Science*, 292–302. Springer.
- Jha, S.; Gulwani, S.; Seshia, S. A.; and Tiwari, A. 2010. Oracle-guided component-based program synthesis. In *ICSE (1)*, 215–224. ACM.
- Jha, S.; and Seshia, S. A. 2017. A theory of formal synthesis via inductive learning. *Acta Informatica*, 54(7): 693–726.
- Kaliszyk, C.; Urban, J.; Michalewski, H.; and Olsák, M. 2018. Reinforcement Learning of Theorem Proving. In *NeurIPS*, 8836–8847.
- Kamil, S.; Cheung, A.; Itzhaky, S.; and Solar-Lezama, A. 2016. Verified lifting of stencil computations. In *PLDI*, 711–726. ACM.
- Kocsis, L.; and Szepesvári, C. 2006. Bandit Based Monte-Carlo Planning. In *ECML*, volume 4212 of *Lecture Notes in Computer Science*, 282–293. Springer.
- Kutsia, T. 2013. Anti-Unification: Algorithms and Applications. In *UNIF@RTA/TLCA*, volume 19 of *EPiC Series in Computing*, 2. EasyChair.
- Laurent, J.; and Platzer, A. 2022. Learning to Find Proofs and Theorems by Learning to Refine Search Strategies: The Case of Loop Invariant Synthesis. In *NeurIPS*.
- Lee, W.; Heo, K.; Alur, R.; and Naik, M. 2018. Accelerating search-based program synthesis using learned probabilistic models. In *PLDI*, 436–449. ACM.

- Martelli, A.; and Montanari, U. 1982. An Efficient Unification Algorithm. *ACM Trans. Program. Lang. Syst.*, 4(2): 258–282.
- Medeiros, L. C.; Aleixo, D. S.; and Lelis, L. H. S. 2022. What Can We Learn Even from the Weakest? Learning Sketches for Programmatic Strategies. In *AAAI*, 7761–7769. AAAI Press.
- Morton, K.; Hallahan, W. T.; Shum, E.; Piskac, R.; and Santolucito, M. 2020. Grammar Filtering for Syntax-Guided Synthesis. In *AAAI*, 1611–1618. AAAI Press.
- Odena, A.; Shi, K.; Bieber, D.; Singh, R.; Sutton, C.; and Dai, H. 2021. BUSTLE: Bottom-Up Program Synthesis Through Learning-Guided Exploration. In *ICLR*. OpenReview.net.
- Plotkin, G. D. 1970. A Note on Inductive Generalization. *Machine Intelligence*, 5: 153–163.
- Raza, M.; Gulwani, S.; and Milic-Frayling, N. 2014. Programming by Example Using Least General Generalizations. In *AAAI*, 283–290. AAAI Press.
- Reynolds, A.; Barbosa, H.; Nötzli, A.; Barrett, C. W.; and Tinelli, C. 2019. cvc4sy: Smart and Fast Term Enumeration for Syntax-Guided Synthesis. In *CAV (2)*, volume 11562 of *Lecture Notes in Computer Science*, 74–83. Springer.
- Reynolds, A.; and Tinelli, C. 2017. SyGuS Techniques in the Core of an SMT Solver. In *SYNT@CAV*, volume 260 of *EPTCS*, 81–96.
- Si, X.; Naik, A.; Dai, H.; Naik, M.; and Song, L. 2020. Code2Inv: A Deep Learning Framework for Program Verification. In *CAV (2)*, volume 12225 of *Lecture Notes in Computer Science*, 151–164. Springer.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T.; Simonyan, K.; and Hassabis, D. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419): 1140–1144.
- Solar-Lezama, A. 2009. The sketching approach to program synthesis. In *Asian Symposium on Programming Languages and Systems*, 4–13. Springer.
- Toni, G. D.; Erculiani, L.; and Passerini, A. 2021. Learning compositional programs with arguments and sampling. *CoRR*, abs/2109.00619.
- Weber, T.; Conchon, S.; Déharbe, D.; Heizmann, M.; Niemetz, A.; and Reger, G. 2019. The SMT Competition 2015-2018. *J. Satisf. Boolean Model. Comput.*, 11(1): 221–259.
- Yernaux, G.; and Vanhoof, W. 2019. Anti-unification in Constraint Logic Programming. *Theory Pract. Log. Program.*, 19(5-6): 773–789.