
Learning to Infer Program Sketches

Maxwell Nye^{1,2} Luke Hewitt^{1,2,3} Joshua Tenenbaum^{1,2,4} Armando Solar-Lezama²

Abstract

Our goal is to build systems which write code automatically from the kinds of specifications humans can most easily provide, such as examples and natural language instruction. The key idea of this work is that a flexible combination of pattern recognition and explicit reasoning can be used to solve these complex programming problems. We propose a method for dynamically integrating these types of information. Our novel intermediate representation and training algorithm allow a program synthesis system to learn, without direct supervision, when to rely on pattern recognition and when to perform symbolic search. Our model matches the memorization and generalization performance of neural synthesis and symbolic search, respectively, and achieves state-of-the-art performance on a dataset of simple English description-to-code programming problems.

1. Introduction

An open challenge in AI is to automatically write code from the kinds of specifications humans can easily provide, such as examples or natural language instruction. Such a system must determine both what the task is and how to write the correct code. Consider writing a simple function which maps inputs to outputs:

$$\begin{aligned} [2, 3, 4, 5, 6] &\rightarrow [2, 4, 6] \\ [5, 8, 3, 2, 2, 1, 12] &\rightarrow [8, 2, 2, 12] \end{aligned}$$

A novice programmer would not recognize from experience any of the program, and would have to *reason* about the entire program structure from first principles. This reasoning would be done by considering the definitions and syntax of the primitives in the programming language, and finding a

way to combine these language constructs to construct an expression with the desired behavior.

A moderately experienced programmer might immediately *recognize*, from learned experience, that because the output list is always a subset of the input list, a `filter` function is appropriate:

```
filter(input, <HOLE>)
```

where `<HOLE>` is a lambda function which filters elements in the list. The programmer would then have to reason about the correct code for `<HOLE>`.

Finally, a programmer very familiar with this domain might immediately recognize both the need for a `filter` function, as well as the correct semantics for the lambda function, allowing the entire program to be written in one shot:

```
filter(input, lambda x: x%2==0)
```

Depending on the familiarity of the domain and the complexity of the problem, humans use a flexible combination of recognition of learned patterns and explicit reasoning to solve programming problems (Lake et al., 2017). Familiar patterns are used, when they exist, and for unfamiliar code elements, explicit reasoning is employed.

This flexibility is not unique to input-output examples. For example, a natural language specification could be used to further specify the desired program, i.e., “Find the even values in a list.” In this case, the process of writing code is analogous. For example, a programmer might learn that “find X in a list” means `filter`, and “even” corresponds to the code `x%2==0`. For a less familiar task, such as “Find values in the list which are powers of two,” a programmer might recognize the need for `filter`, but would need to reason about how to write a lambda function which classifies powers of two.

We propose a system which mimics the human ability to dynamically incorporate pattern recognition and reasoning to solve programming problems from examples or natural language specification. We show that without direct supervision, our model learns to find good intermediates between pattern recognition and symbolic reasoning components, and outperforms existing models on several programming tasks.

¹MIT Brain and Cognitive Sciences ²MIT CSAIL ³MIT-IBM AI Lab ⁴Center for Brains, Minds and Machines (CBMM). Correspondence to: Maxwell Nye <mnye@mit.edu>.

Recent work (Murali et al., 2017; Dong & Lapata, 2018) has attempted to combine learned pattern recognition and explicit reasoning using *program sketches*—schematic outlines of full programs (Solar-Lezama, 2008). In Murali et al. (2017), a neural network is trained to output program sketches when conditioned on a spec, and candidate sketches are converted into full programs using symbolic synthesis techniques, which approximate explicit reasoning from first principles.

However, previous systems use static, hand-designed intermediate sketch grammars, which do not allow the system to learn how much to rely on pattern recognition and how much to rely on symbolic search. The neural network is trained to map from spec to a pre-specified sketch, and cannot learn to output a more detailed sketch, if the pattern matching task is easy, or learn to output a more general sketch, if the task is too difficult.

Ideally, a neuro-symbolic synthesis system should dynamically take advantage of the relative strengths of its components. When given an easy or familiar programming task, for example, it should rely on its learned pattern recognition, and output a fuller program with a neural network, so that less time is required for synthesis. In contrast, when given a hard task, the system should learn to output a less complete sketch and spend more time filling in the sketch with search techniques. We believe that this flexible integration of neural and symbolic computation, inspired by humans, is necessary for powerful, domain-general intelligence, and for solving difficult programming tasks.

The key idea in this work is to allow a system to *learn* a suitable intermediate sketch representation between a learned neural proposer and a symbolic search mechanism. Inspired by Murali et al. (2017), our technique comprises a learned **neural sketch generator** and an enumerative symbolic **program synthesizer**. In contrast to previous work, however, we use a flexible and domain-general sketch grammar, and a novel self-supervised training objective, which allows the network to learn how much to rely on each component of the system. The result is a flexible, domain-general program synthesis system, which has the ability to learn sophisticated patterns from data, comparably to Devlin et al. (2017), as well as utilize explicit symbolic search for difficult or out-of-sample problems, as in Balog et al. (2016).

Without explicit supervision, our model learns good intermediates between neural network and synthesis components. This allows our model to increase data efficiency and generalize better to out-of-sample test tasks compared to RNN-based models. Concretely, our contributions are as follows:

- We develop a novel neuro-symbolic program synthesis system, which writes programs from input-output examples and natural language specification by learning

a suitable intermediate sketch representation between a neural network sketch generator and a symbolic synthesizer.

- We introduce a novel training objective, which we used to train our system to find suitable sketch representations without explicit supervision.
- We validate our system by demonstrating our results in two programming-by-example domains, list processing problems and string transformation problems, and achieve state-of-the-art performance on the *AlgoLisp English-to-code* test dataset.

2. Problem Formulation

Assume that we have a DSL which defines a space of programs, \mathcal{G} . In addition, we have a set of program specifications, or *specs*, which we wish to ‘solve’. We assume each spec \mathcal{X}_i is satisfied by some true unknown program F_i .

If our specification contains a set of IO examples $\mathcal{X}_i = \{(x_{ij}, y_{ij})\}_{j=1..n}$, then we can say that we have solved a task \mathcal{X}_i if we find the true program F_i , which must satisfy all of the examples:

$$\forall j : F_i(x_{ij}) = y_{ij}$$

Our goal is to build a system which, given \mathcal{X}_i , can quickly recover F_i . For our purposes, *quickly* is taken to mean that such a solution is found within some threshold time, $\text{Time}(\mathcal{X}_i \rightarrow F_i) < t$. Formally, then, we wish to maximize the probability that our system solves each problem within this threshold time:

$$\max \log \mathbb{P} \left[\text{Time}(\mathcal{X}_i \rightarrow F_i) < t \right] \quad (1)$$

Additionally, for some domains our spec \mathcal{X} may contain additional informal information, such as natural language instruction. In this case, we can apply same formalism, maximizing the probability that the true program F_i is found given the spec \mathcal{X}_i , within the threshold time.

3. Our Approach: Learning to Infer Sketches

3.1. System Overview:

Our approach, inspired by work such as Murali et al. (2017), is to represent the relationship between program specification and program using an intermediate representation called a program sketch. However in contrast to previous work, where the division of labor between generating sketches and filling them in is fixed, our approach allows this division of labor to be learned, without additional supervision. We define a *sketch* simply as a valid program tree in the DSL, where any number of subtrees has been replaced by

a special token: `<HOLE>`. Intuitively, this token designates locations in the program tree for which pattern-based recognition is difficult, and more explicit search methods are necessary.

Our system consists of two main components: 1) a **sketch generator**, and 2) a **program synthesizer**.

The **sketch generator** is a distribution over program sketches given the spec: $q_\phi(\text{sketch}|\mathcal{X})$. The generator is parametrized by a recurrent neural network, and is trained to assign high probability to sketches which are likely to quickly yield programs satisfying the spec when given to the synthesizer. Details about the learning scheme and architecture will be discussed below.

The **program synthesizer** takes a sketch as a starting point, and performs an explicit symbolic search to “fill in the holes” in order to find a program which satisfies the spec.

Given a set of test problems, in the form of a set of specs, the system searches for correct programs as follows: The sketch generator, conditioned on the program spec, outputs a distribution over sketches. A fixed number of candidate sketches $\{s_i\}$ are extracted from the generator. This set $\{s_i\}$ is then given to the program synthesizer, which searches for full programs maximizing the likelihood of the spec. For each candidate sketch, the synthesizer uses symbolic enumeration techniques to search for full candidate programs which are formed by filling in the holes in the sketch.

Using our approach, our system is able to flexibly learn the optimal amount of detail needed in the sketches, essentially learning how much to rely on each component of the system. Furthermore, due to our domain-general sketch grammar, we are easily able to implement our system in multiple different domains with very little overhead.

3.2. Learning to Infer Sketches via Self-supervision

By using sketches as an intermediate representation, we reframe our program synthesis problem (Eq. 1) as follows: learn a sketch generator $q_\phi(s|\mathcal{X})$ which, given a spec \mathcal{X}_i , produces a ‘good’ sketch s from which the synthesizer can quickly find the solution F_i . We may thus wish to maximize the probability that our sketch generator produces a ‘good’ sketch:

$$\max_{\phi} \log \mathbb{P}_{s \sim q_\phi(-|\mathcal{X}_i)} [\text{Time}(s \rightarrow F_i) < t] \quad (2)$$

where $\text{Time}(s \rightarrow F_i)$ is the time taken for the synthesizer to discover the solution to \mathcal{X}_i by filling the holes in sketch s , and t is the synthesizer’s evaluation budget.

In order to learn a system which is most robust, we make one final modification to Equation (2): at train time we do not necessarily know what the timeout will be during evaluation, so we would like to train a system which is agnostic to the

amount of time it would have. Ideally, if a program can be found entirely (or almost entirely) using familiar patterns, then the sketch generator should assign high probability to very complete sketches. However, if a program is unfamiliar or difficult, the sketches it favors should be very general, so that the synthesizer must do most of the computation. To do this, we can train the generator to output sketches which would be suitable for a wide distribution of evaluation budgets. This can be achieved by allowing the budget t to be a random variable, sampled from some distribution \mathcal{D}_t . Adding this uncertainty to Equation (2) yields:

$$\max_{\phi} \log \mathbb{P}_{\substack{t \sim \mathcal{D}_t \\ s \sim q_\phi(-|\mathcal{X}_i)}} [\text{Time}(s \rightarrow F_i) < t] \quad (3)$$

In practice, we can achieve this maximization by self-supervised training. That is, given a dataset of program-spec pairs, for each spec we optimize the *generator* to produce only the sketches from which we can quickly recover its underlying program. Thus, given training data as (F, \mathcal{X}) pairs, our training objective may be written:

$$\text{obj} = \mathbb{E}_{\substack{t \sim \mathcal{D}_t \\ (F, \mathcal{X}) \sim \mathcal{G}}} \log \sum_{s: \text{Time}(s \rightarrow F) < t} q_\phi(s|\mathcal{X}) \quad (4)$$

During each step of training, t is sampled from \mathcal{D}_t , and the network is trained to assign high probability to those sketches which can be synthesized into a full program within the budget t . Using this training scheme, the network learns to output a distribution of sketches, some of which are very specific and can be synthesized quickly, while others are more general but require more time to synthesize into full programs. This allows the system to perform well with various enumeration budgets and levels of problem difficulty: the system quickly solves “easy” problems with very “concrete” sketches, but also samples more general sketches, which can be used to solve difficult problems for which the system’s learned inductive biases are less appropriate.

4. Our Implementation

In this section, we discuss our implementation of the above ideas which we use to solve the list processing and string editing tasks discussed above, in a system we call SKETCHADAPT.

4.1. Seq-to-Seq Neural Sketch Generator

For our sketch generator, we use a sequence-to-sequence recurrent neural network with attention, inspired by Devlin et al. (2017) and Bunel et al. (2018). Our model is inspired by the “Att-A” model in Devlin et al. (2017): the model encodes the spec via LSTM encoders, and then decodes a program token-by-token while attending to the spec. To facilitate the learning of the output grammar, our model




Figure 1. Schematic overview of our model. A program spec (in the form of examples) is fed into a sketch generator, which outputs a distribution over sketches. In our experiments, the neural sketch generator is parametrized by a seq-to-seq RNN with attention. The program sketch is given to a program synthesizer, which searches for full programs which satisfy the spec. Our enumerative synthesizer is guided by a learned recognizer, which is conditioned on the spec and the sketch and predicts the likelihood of using each program token to fill in the sketch.

also has an additional learned LSTM language model as in Bunel et al. (2018), which reweights the program token probabilities from the seq-to-seq model. This LSTM language model is simply trained to assign high probability to likely sequences of tokens via supervised learning.

4.2. Synthesis via Enumeration

Our symbolic sketch synthesizer is based on Ellis et al. (2018) and Balog et al. (2016) and has two components: a breadth-first probabilistic enumerator, which enumerates candidate programs from most to least likely, and a neural recognizer, which uses the spec to guide this probabilistic enumeration.

The enumerator, based on Ellis et al. (2018) uses a strongly typed DSL, and assumes that all programs are expressions in λ -calculus. Each primitive has an associated production probability. These production probabilities constitute the parameters, θ , for a probabilistic context free grammar, thus inducing a distribution over programs $p(F|s, \theta)$. Synthesis proceeds by enumerating candidate programs which satisfy a sketch in decreasing probability under this distribution. Enumeration is done in parallel for all the candidate sketches, until a full program is found which satisfies the input-output examples in the spec, or until the enumeration budget is exceeded.

The learned recognizer is inspired by Menon et al. (2013) and the “Deepcoder” system in Balog et al. (2016). For a given task, an RNN encodes each spec into a latent vector. The latent vectors are averaged, and the result is passed into a feed-forward MLP, which terminates in a softmax layer. The resulting vector is used as the set of production probabilities θ which the enumerator uses to guide search.

SKETCHADAPT succeeds by exploiting the fundamental difference in search capabilities between its neural and symbolic components. Pure-synthesis approaches can enumerate and check candidate programs extremely quickly—we enumerate roughly 3×10^3 prog/sec, and the fastest enu-

merator for list processing exceeds 10^6 prog/sec. However, generating expressions larger than a few nodes requires searching an exponentially large space, making enumeration impractical for large programs. Conversely, seq2seq networks (and tree RNNs) require fewer samples to find a solution but take much more time per sample (many milliseconds per candidate in a beam search) so are restricted to exploring only hundreds of candidates. They therefore succeed when the solution is highly predictable (even if it is long), but fail if even a small portion of the program is too difficult to infer. By flexibly combining these two approaches, our system searches the space of programs more effectively than either approach alone; SKETCHADAPT uses learned patterns to guide a beam search when possible, and fast enumeration for portions of the program which are difficult to recognize. This contrasts with Murali et al. (2017), where the division of labor is fixed and cannot be learned.

4.3. Training

The training objective above (Eq. 4) requires that for each training program F , we know the set of sketches which can be synthesized into F in less than time t (where the synthesis time is given by $\text{Time}(s \rightarrow F)$.) A simple way to determine this set would be to simulate synthesis for each candidate sketch, to determine synthesis can succeed in less time than t . In practice, we do not run synthesis during training of the sketch generator to determine $\text{Time}(s \rightarrow F)$. One benefit of the probabilistic enumeration is that it provides an estimate of the enumeration time of a sketch. It is easy to calculate the likelihood $p(F|s, \theta)$ of any full program F , given sketch s and production probabilities θ given by the recognizer $\theta = r(\mathcal{X})$. Because we enumerate programs in decreasing order of likelihood, we know that search time (expressed as number of evaluations) can be upper bounded using the likelihood: $\text{Time}(s \rightarrow F) \leq [p(F|s, \theta)]^{-1}$. Thus, we can use the inverse likelihood to lower bound Equation (4) by:

$$obj \geq \mathbb{E}_{\substack{t \sim \mathcal{D}_t \\ (F, \mathcal{X}) \sim \mathcal{G}}} \log \sum_{s: p^{-1}(F|s, \theta) < t} q_\phi(s|\mathcal{X}) \quad (5)$$

While it is often tractable to evaluate this sum exactly, we may further reduce computational cost if we can identify a smaller set sketches which dominate the log sum. Fortunately we observe that the generator and synthesizer are likely to be highly correlated, as each program token must be explained by either one or the other. That is, sketches which maximize $q_\phi(s|\mathcal{X})$ will typically minimize $p(F|s, \theta)$. Therefore, we might hope to find a close bound on Equation (5) by summing only the few sketches that minimize $p(F|s, \theta)$. In this work we have found it sufficient to use only a *single minimal sketch*, yielding the objective obj^* :

$$obj^* = \mathbb{E}_{\substack{t \sim \mathcal{D}_t \\ (F, \mathcal{X}) \sim \mathcal{G}}} \log q_\phi(s^*|\mathcal{X}) \leq obj,$$

$$\text{where } s^* = \arg \min_{s: p^{-1}(F|s, \theta) < t} p(F|s, \theta) \quad (6)$$

This allows us to perform a much simpler and more practical training procedure, maximizing a lower bound of our desired objective. For each full program sampled from the DSL, we sample a timeout $t \sim \mathcal{D}_t$, and determine the sketch with maximum likelihood, for which $p^{-1}(F|s, \theta) < t$. We then train the neural network to maximize the probability of that sketch. Intuitively, we are sampling a random timeout, and training the network to output the easiest sketch which still solves the task within that timeout.

For each full program F , we assume that the set of sketches which can be synthesized into F will have enumeration times distributed roughly exponentially. Therefore, in order for our training procedure to utilize a range of sketch sizes, we use an exponential distribution for the timeout: $t \sim \text{Exp}(\alpha)$, which works well in practice.

Our training methodology is described in Algorithm 1, and the evaluation approach is described in Algorithm 2.

5. Experiments

We provide the results of evaluating SKETCHADAPT in three test domains. For all test domains, we compare against two alternate models, which can be regarded as lesioned versions of our model, as well as existing models in the literature:

The “**Synthesizer only**” alternate model is equivalent to our program synthesizer module, using a learned recognition model and enumerator. Instead of enumerating from holes in partially filled-in sketches, the “Synthesizer only” model enumerates all programs from scratch, starting from a single **<HOLE>** token. This model is comparable to the “Deepcoder” system in Balog et al. (2016), which was developed to solve the list transformation tasks we examine in subsection 5.1.

The “**Generator only**” alternate model is a fully seq-to-seq RNN, equivalent in architecture to our sketch generator, trained simply to predict the entire program. This model

Algorithm 1 SKETCHADAPT Training

Require: Sketch Generator $q_\phi(\text{sketch}|\mathcal{X})$; Recognizer $r_\psi(\mathcal{X}, \text{sketch})$; Enumerator dist. $p(F|\theta, \text{sketch})$, Base Parameters θ_{base}

Train Recognizer, r_ψ :

for F, \mathcal{X} in Dataset (or sampled from DSL) **do**

 Sample $t \sim \mathcal{D}_t$

sketches, probs \leftarrow list all possible sketches of F ,
 with probs given by $p(F|s, \theta_{base})$

sketch \leftarrow sketch with largest prob s.t. prob $< t$.

$\theta \leftarrow r_\psi(\mathcal{X}, \text{sketch})$

 grad. step on ψ to maximize $\log p(F|\theta, \text{sketch})$

end for

Train Sketch Generator, q_ϕ :

for F, \mathcal{X} in Dataset (or sampled from DSL) **do**

 Sample $t \sim \mathcal{D}_t$

$\theta \leftarrow r_\psi(\mathcal{X})$

sketches, probs \leftarrow list all possible sketches of F ,
 with probs given by $p(F|s, \theta)$

sketch \leftarrow sketch with largest prob s.t. prob $< t$.

 grad. step on ϕ to maximize $\log q_\phi(\text{sketch}|\mathcal{X})$

end for

is comparable to the “RobustFill” model in Devlin et al. (2017), which was developed to solve the string transformation tasks we examine in subsection 5.2. This model is also comparable to the sequence-to-sequence models in Polosukhin & Skidanov (2018).

5.1. List Processing

In our first, small scale experiment, we examine problems that require an agent to synthesize programs which transform lists of integers. We use the list processing DSL from Balog et al. (2016), which consists of 34 unique primitives. The primitives consist of first order list functions, such as

Algorithm 2 SKETCHADAPT Evaluation

Require: Sketch Generator $q_\phi(\text{sketch}|\mathcal{X})$; Recognizer $r_\psi(\mathcal{X}, \text{sketch})$; Enumerator dist. $p(F|\theta, \text{sketch})$

function synthesizeProgram(\mathcal{X})

sketches \leftarrow beam search $q_\phi(\cdot|\mathcal{X})$

for *sketch* in *sketches* {in parallel} **do**

$\theta_{sketch} \leftarrow r_\psi(\mathcal{X}, \text{sketch})$

while timeout not exceeded **do**

$F \leftarrow$ next full prog. from enumerate(*sketch*, θ_{sketch})

if F satisfies \mathcal{X} **then**

return F

end if

end while

end for

Table 1. Example list processing programs

Input	Output	Program
1, [-101, 63, 64, 79, 119, 91, -56, 47, -74, -33]	39	(MAXIMUM (MAP DIV3 (DROP input0 input1)))
4, [-6, -96, -45, 17, 26, -38, 17, -18, -112, -48]	8	
2, [-9, 5, -8, -9, 9, -3, 7, -5, -10, 1]	[100, 16]	(TAKE input0 (MAP SQR (MAP DEC input1)))
3, [-5, -8, 0, 10, 2, -7, -3, -5, 6, 2]	[36, 81, 1]	




Figure 2. Left: Results of model trained on list processing programs of length 3, using a beam size of 100, plotted as a function of the enumeration budget. Right: Generalization results: models trained on list processing programs of length 3, evaluated on programs of length 4. Although the synthesis-only Deepcoder model was developed for the list processing problems, our SKETCHADAPT model requires much less enumeration to achieve high accuracy for the within-sample length 3 programs, and performs comparably for the out-of-sample length 4 programs, far exceeding the “Generator only” RNN-based model.




Figure 3. Performance on string editing problems. Although RobustFill was developed for string editing problems, SKETCHADAPT achieves higher accuracy on these tasks.

head, last and reverse, higher order list functions, such as map, filter and zipwith, and lambda functions such as min, max and negate. Our programs are semantically equivalent, but differ from those in Balog et al. (2016) in that we use no bound variables (apart from input variables), and instead synthesize a single s-expression.

As in Balog et al. (2016), the spec for each task is a small number of input-output examples. See Table 1 for sample programs and examples.

Our goal was to determine how well our system could perform in two regimes: *within-sample*, where the test data is similar to the training data, and *out-of-sample*, where the test data distribution is different from the training distribution. We trained our model on programs of length 3, and tested its performance two datasets, one consisting of 100 programs of length 3, and the other with 100 length 4 programs. With these experiments, we could determine how well our system synthesizes easy and familiar programs (length 3), and difficult programs which require generalization (length 4).

During both training and evaluation, the models were conditioned on 5 example input-output pairs, which contain integers with magnitudes up to 128. In Figure 2, we plot the proportion of tasks solved as a function of the number of candidate programs enumerated per task.

Although a “Generator only” RNN model is able to synthesize many length 3 programs, it performs very poorly on the out-of-sample length 4 programs. We also observe that, while the “Synthesizer only” model can take advantage of a large enumeration budget and solve a higher propor-

Table 2. Example string editing programs

Input	Output	Program
'Madelaine'	'M-'	(concat_list (expr GetUpTo_Char) (concat_single (Constant dash)))
'Olague'	'O-'	
'118-980-214'	'214'	(concat_single (expr GetToken_Number_-1))
'938-242-504'	'504'	

tion of out-of-sample tasks than the “Generator only” RNN, it does not take advantage of learned patterns to synthesize the length 3 programs quickly, due to poor inductive biases. Only our model is able to perform well on both within-sample and out-of-sample tasks.

5.2. String Transformations

In our second test domain, we explored programs which perform string transformations, as in Gulwani (2011). These problems involve finding a program which maps an input string to an output string. Typically, these programs are used to manipulate the syntactic form of the underlying information, with minimal changes to the underlying semantics. Examples include converting a list of 'FirstName LastName' to 'LastInitial, Firstname'. These problems have been studied by Gulwani (2011); Polozov & Gulwani (2015); Devlin et al. (2017) and others. We show that our system is able to accurately recover these programs.

As our test corpus, we used string editing problems from the SyGuS (Alur et al., 2016) program synthesis competition, and string editing tasks used in Ellis et al. (2018). We excluded tasks requiring multiple input strings or a pre-specified string constant, leaving 48 SyGuS programs and 79 programs from Ellis et al. (2018). Because we had a limited corpus of problems, we trained our system on synthetic data only, sampling all training programs from the DSL.

Because our system has no access to the test distribution, this domain allows us to see how well our method is able to solve real-world problems when trained only on a synthetic distribution.

Furthermore, the string editing DSL is much larger than the list transformation DSL. This means that enumerative search is both slower and less effective than for the list transformation programs, where a fast enumerator could brute force the entire search space (Balog et al., 2016). Because of this, the “Synthesizer only” model is not able to consistently enumerate sufficiently complex programs from scratch.

We trained our model using self-supervision, sampling training programs randomly from the DSL and conditioning the models on 4 examples of input-output pairs, and evaluated on our test corpus. We plot our results in Figure 3.

Overall, SKETCHADAPT outperforms the “Synthesizer




Figure 4. Algolisp: varying training data size. We trained our model and baselines on various dataset sizes, and evaluated performance on a held-out test dataset. Our SKETCHADAPT system considerably outperforms baselines in the low-data regime. “Generator only” model, and matches or exceeds the performance of the “Generator only” RNN model, which is noteworthy given that it is equivalent to RobustFill, which was designed to synthesize string editing programs. We also note that the beam size used in evaluation of the “Generator only” RNN model has a large impact on performance. However, the performance of our SKETCHADAPT system is less dependent on beam size, suggesting that the system is able to effectively supplement a smaller beam size with enumeration.

5.3. Algolisp: Description to Programs

Our final evaluation domain is the Algolisp DSL and dataset, introduced in Polosukhin & Skidanov (2018). The Algolisp dataset consists of programs which manipulate lists of integers and lists of strings. In addition to input-output examples, each specification also contains a natural language description of the desired program. We use this dataset to examine how well our system can take advantage of highly unstructured specification data such as natural language, in addition to input-output examples.

The Algolisp problems are very difficult to solve using only examples, due to the very large search space and program complexity (see Table 4: SKETCHADAPT, IO only). However, the natural language description makes it possible, with enough data, to learn a highly accurate semantic parsing scheme and solve many of the test tasks. In addition, because this domain uses real data, and not data generated from self-supervision, we wish to determine how data-efficient our algorithm is. Therefore, we train our model on

Table 3. Example problems from the AlgoLisp dataset

Spec	Program
Consider an array of numbers, find elements in the given array not divisible by two	<pre>(filter a (lambda1 (== (% arg1 2) 1)))</pre>
You are given an array of numbers, your task is to compute median in the given array with its digits reversed	<pre>(reduce (reverse (digits (deref (sort a) (/ (len a) 2)))) 0 (lambda2 (+ (* arg1 10) arg2)))</pre>

Table 4. Algolisp results on full dataset

Model	Full dataset		Filtered ¹	
	(Dev)	Test	(Dev)	Test
SKETCHADAPT (Ours)	(88.8)	90.0	(95.0)	95.8
Synthesizer only	(5.2)	7.3	(5.6)	8.0
Generator only	(91.4)	88.6	(98.4)	95.6
SKETCHADAPT, IO only	(4.9)	8.3	(5.6)	8.8
Seq2Tree+Search	(86.1)	85.8	-	-
SAPS ²	(83.0)	85.2	(93.2)	92.0

Table 5. Algolisp generalization results: Trained on 8000 programs, excluding ‘Odd’ concept:

Model	Even	Odd
SKETCHADAPT (Ours)	34.4	29.8
Synthesizer only	23.7	0.0
Generator only	4.5	1.1

subsets of the data of various sizes to test generalization.

Figure 4 and Table 4 depict our main results for this domain, testing all systems with a maximum timeout of 300 seconds per task.¹ When using a beam size of 10 on the full dataset, SKETCHADAPT and the ‘‘Generator only’’ RNN baseline far exceed previously reported state of art performance and achieve near-perfect accuracy, whereas the ‘‘Synthesizers only’’ model is unable to achieve high performance. However, when a smaller number of training programs is used, SKETCHADAPT significantly outperforms the ‘‘Generator only’’ RNN baseline. These results indicates that the symbolic search allows SKETCHADAPT to perform stronger generalization than pure neural search methods.

Strong generalization to unseen subexpressions: As a final test of generalization, we trained SKETCHADAPT and our baseline models on a random sample of 8000 training programs, excluding all those which contain the function ‘odd’ as expressed by the AlgoLisp subexpression `(lambda1(== (% arg1 2) 1))` (in python, `lambda x: x%2==1`). We then evaluate on all 635 test programs containing ‘odd’, as well as the 638 containing ‘even’ (`lambda x: x%2==0`). As shown in Table 5, the ‘‘Generator only’’ RNN baseline exhibits only weak generalization, solving novel tasks which require the ‘even’ subex-

¹As in Bednarek et al. (2018), we filter the test and dev datasets for only those tasks for which reference programs satisfy the given specs. The ‘‘filtered’’ version is also used for Figure 4.

²Introduced in Bednarek et al. (2018)

pression but not those which require the previously unseen ‘odd’ subexpression. By contrast, SKETCHADAPT exhibits strong generalization to both ‘even’ and ‘odd’ programs.

6. Related Work

Our work takes inspiration from the neural program synthesis work of Balog et al. (2016), Devlin et al. (2017) and Murali et al. (2017). Much recent work has focused on learning programs using deep learning, as in Kalyan et al. (2018), Bunel et al. (2018), Shin et al. (2018), or combining symbolic and learned components, such as Parisotto et al. (2016), Kalyan et al. (2018), Chen et al. (2017), Zohar & Wolf (2018), and Zhang et al. (2018). Sketches have also been explored for semantic parsing (Dong & Lapata, 2018) and differentiable programming (Bošnjak et al., 2017). We also take inspiration from the programming languages literature, particularly Sketch (Solar-Lezama, 2008) and angelic nondeterminism (Bodik et al., 2010). Other work exploring symbolic synthesis methods includes λ^2 (Feser et al., 2015) and Schkufza et al. (2016). Learning programs has also been studied from a Bayesian perspective, as in EC (Dechter et al., 2013), Bayesian Program Learning (Lake et al., 2015), and inference compilation (Le et al., 2016).

7. Discussion

We developed a novel neuro-symbolic scheme for synthesizing programs from examples and natural language. Our system, SKETCHADAPT, combines neural networks and symbolic synthesis by learning an intermediate ‘sketch’ representation, which dynamically adapts its specificity for each task. Empirical results show that SKETCHADAPT recognizes common motifs as effectively as pure RNN approaches, while matching or exceeding the generalization of symbolic synthesis methods. We believe that difficult program synthesis tasks cannot be solved without flexible integration of pattern recognition and explicit reasoning, and this work provides an important step towards this goal.

We also hypothesize that learned integration of different forms of computation is necessary not only for writing code, but also for other complex AI tasks, such as high-level planning, rapid language learning, and sophisticated question answering. In future work, we plan to explore the ideas presented here for other difficult AI domains.

Acknowledgements

The authors would like to thank Kevin Ellis and Lucas Morales for very useful feedback, as well as assistance using the [EC codebase](#). M. N. is supported by an NSF Graduate Fellowship and an MIT BCS Hilibrand Graduate Fellowship. L. H. is supported by the MIT-IBM Watson AI Lab.

References

- Alur, R., Fisman, D., Singh, R., and Solar-Lezama, A. Sygus-comp 2016: Results and analysis. *arXiv preprint arXiv:1611.07627*, 2016.
- Balog, M., Gaunt, A. L., Brockschmidt, M., Nowozin, S., and Tarlow, D. Deepcoder: Learning to write programs. *arXiv preprint arXiv:1611.01989*, 2016.
- Bednarek, J., Piaskowski, K., and Krawiec, K. Ain’t nobody got time for coding: Structure-aware program synthesis from natural language. *arXiv preprint arXiv:1810.09717*, 2018.
- Bodik, R., Chandra, S., Galenson, J., Kimelman, D., Tung, N., Barman, S., and Rodarmor, C. Programming with angelic nondeterminism. In *ACM Sigplan Notices*, volume 45, pp. 339–352. ACM, 2010.
- Bošnjak, M., Rocktäschel, T., Naradowsky, J., and Riedel, S. Programming with a differentiable forth interpreter. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 547–556. JMLR. org, 2017.
- Bunel, R., Hausknecht, M., Devlin, J., Singh, R., and Kohli, P. Leveraging grammar and reinforcement learning for neural program synthesis. *arXiv preprint arXiv:1805.04276*, 2018.
- Chen, X., Liu, C., and Song, D. Towards synthesizing complex programs from input-output examples. *arXiv preprint arXiv:1706.01284*, 2017.
- Dechter, E., Malmaud, J., Adams, R. P., and Tenenbaum, J. B. Bootstrap learning via modular concept discovery. In *IJCAI*, pp. 1302–1309, 2013.
- Devlin, J., Uesato, J., Bhupatiraju, S., Singh, R., Mohamed, A.-r., and Kohli, P. Robustfill: Neural program learning under noisy i/o. *arXiv preprint arXiv:1703.07469*, 2017.
- Dong, L. and Lapata, M. Coarse-to-fine decoding for neural semantic parsing. *arXiv preprint arXiv:1805.04793*, 2018.
- Ellis, K., Morales, L., Sablé-Meyer, M., Solar-Lezama, A., and Tenenbaum, J. Learning libraries of subroutines for neurally-guided bayesian program learning. *NIPS*, 2018.
- Feser, J. K., Chaudhuri, S., and Dillig, I. Synthesizing data structure transformations from input-output examples. In *ACM SIGPLAN Notices*, volume 50, pp. 229–239. ACM, 2015.
- Gulwani, S. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, pp. 317–330. ACM, 2011.
- Kalyan, A., Mohta, A., Polozov, O., Batra, D., Jain, P., and Gulwani, S. Neural-guided deductive search for real-time program synthesis from examples. 2018.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- Le, T. A., Baydin, A. G., and Wood, F. Inference compilation and universal probabilistic programming. *arXiv preprint arXiv:1610.09900*, 2016.
- Menon, A., Tamuz, O., Gulwani, S., Lampson, B., and Kalai, A. A machine learning framework for programming by example. In *International Conference on Machine Learning*, pp. 187–195, 2013.
- Murali, V., Qi, L., Chaudhuri, S., and Jermaine, C. Neural sketch learning for conditional program generation. *arXiv preprint arXiv:1703.05698*, 2017.
- Parisotto, E., Mohamed, A.-r., Singh, R., Li, L., Zhou, D., and Kohli, P. Neuro-symbolic program synthesis. *arXiv preprint arXiv:1611.01855*, 2016.
- Polosukhin, I. and Skidanov, A. Neural program search: Solving programming tasks from description and examples. *arXiv preprint arXiv:1802.04335*, 2018.
- Polozov, O. and Gulwani, S. Flashmeta: a framework for inductive program synthesis. In *ACM SIGPLAN Notices*, volume 50, pp. 107–126. ACM, 2015.
- Schkufza, E., Sharma, R., and Aiken, A. Stochastic program optimization. *Communications of the ACM*, 59(2):114–122, 2016.
- Shin, R., Polosukhin, I., and Song, D. Improving neural program synthesis with inferred execution traces. In *Advances in Neural Information Processing Systems*, pp. 8931–8940, 2018.
- Solar-Lezama, A. *Program synthesis by sketching*. PhD thesis, University of California, Berkeley, 2008.

Zhang, L., Rosenblatt, G., Fetaya, E., Liao, R., Byrd, W., Might, M., Urtasun, R., and Zemel, R. Neural guided constraint logic programming for program synthesis. In *Advances in Neural Information Processing Systems*, pp. 1744–1753, 2018.

Zohar, A. and Wolf, L. Automatic program synthesis of long programs with a learned garbage collector. In *Advances in Neural Information Processing Systems*, pp. 2098–2107, 2018.