



HAL
open science

Scaling Neural Program Synthesis with Distribution-Based Search

Nathanaël Fijalkow, Guillaume Lagarde, Théo Matricon, Kevin Ellis, Pierre Ohlmann, Akarsh Nayan Potta

► **To cite this version:**

Nathanaël Fijalkow, Guillaume Lagarde, Théo Matricon, Kevin Ellis, Pierre Ohlmann, et al.. Scaling Neural Program Synthesis with Distribution-Based Search. Proceedings of the AAAI Conference on Artificial Intelligence, 2022, 36 (6), pp.6623-6630. 10.1609/aaai.v36i6.20616 . hal-03775025

HAL Id: hal-03775025

<https://cnrs.hal.science/hal-03775025v1>

Submitted on 26 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Scaling Neural Program Synthesis with Distribution-based Search

Nathanaël Fijalkow^{1,2}, Guillaume Lagarde¹, Théo Matricon¹,
Kevin Ellis³, Pierre Ohlmann⁴, and Akarsh Potta⁵

¹CNRS, LaBRI and Université de Bordeaux, France

²The Alan Turing Institute of data science, United Kingdom

³Cornell University, United States

⁴University of Paris, France

⁵Indian Institute of Technology Bombay, India

Abstract

We consider the problem of automatically constructing computer programs from input-output examples. We investigate how to augment probabilistic and neural program synthesis methods with new search algorithms, proposing a framework called distribution-based search. Within this framework, we introduce two new search algorithms: `HEAP SEARCH`, an enumerative method, and `SQRT SAMPLING`, a probabilistic method. We prove certain optimality guarantees for both methods, show how they integrate with probabilistic and neural techniques, and demonstrate how they can operate at scale across parallel compute environments. Collectively these findings offer theoretical and applied studies of search algorithms for program synthesis that integrate with recent developments in machine-learned program synthesizers.

1 Introduction

Writing software is tedious, error-prone, and accessible only to a small share of the population – yet coding grows increasingly important as the digital world plays larger and larger roles in peoples’ lives. Program synthesis seeks to make coding more reliable and accessible by developing methods for automatically constructing programs Gulwani et al. (2017). For example, the FlashFill system Gulwani et al. (2017) in Microsoft Excel makes coding more accessible by allowing nontechnical users to synthesize spreadsheet programs by giving input-output examples, while the TF-coder system Shi et al. (2020a) seeks to make coding neural networks more reliable by synthesizing TensorFlow code from input-outputs. Where these systems have been most successful is when they pair a specialized *domain-specific language* (DSL) to a domain-specific search algorithm for synthesizing programs in that DSL. A recent trend – both in industry Kalyan et al. (2018) and academia Devlin et al. (2017) – is to employ machine learning methods to learn to quickly search for a program in the DSL Balog et al. (2017); Devlin et al. (2017); Lee et al. (2018); Zhang et al.

(2018); Polosukhin and Skidanov (2018); Kalyan et al. (2018); Zohar and Wolf (2018); Chen et al. (2018). Many such recent works have explored engineering better neural networks for guiding program search, effectively by training the network to act as a language model over source code that conditions on input-outputs Polosukhin and Skidanov (2018). Here, we ‘pop up’ a level and instead ask: given a neural net that probabilistically generates source code, how can we most efficiently deploy that model in order to find a program consistent with some input-outputs? This concern arises because program synthesis requires solving a hard combinatorial search problem (exploring a possibly infinite space of programs), so taming this combinatorial explosion makes the difference between a practically useful system, and a system which cannot scale to anything but the most trivial of programs.

At a high-level the approaches we develop in this work follow a 2-stage pipeline: in the first stage a learned model predicts probabilistic weights, and in the second stage a symbolic search algorithm uses those weights to explore the space of source code. Our contributions target the second stage of this pipeline, and we focus on theoretical analysis of sampling-based search algorithms, new search algorithms based on neurally-informed enumeration, and empirical evaluations showing that recent neural program synthesizers can compose well with our methods.

This 2-stage pipelined approach has several benefits. The first is that the cost of querying the neural network is usually very small compared to the cost of combinatorial search, yet in practice the neural model learns to provide rough-and-ready probabilistic predictions to guide the search. A second benefit is that even if the probabilistic predictions are inaccurate, our methods maintain soundness and completeness (but may take longer to run). Another appeal is that it can be naturally combined with other classical approaches for program synthesis.

Our contributions:

- A theoretical framework called distribution-based search for evaluating and comparing search algorithms in the context of machine-learned predictions.
- Two new search algorithms: HEAP SEARCH, an enumerative method, and SQRT SAMPLING, a probabilistic method. We prove a number of theoretical results about them, in particular that they are both loss optimal.
- A method for running any search algorithm across parallel compute environments.

We perform an empirical evaluation of existing and new search algorithms, showing how the new methods integrate with probabilistic and neural techniques.

2 Distribution-based search

We work within the syntax guided program synthesis (SyGuS) framework introduced by Alur et al. (2013), see also Alur et al. (2018). In this setting, the DSL is given by a set of primitives together with their (possibly polymorphic) types and semantics.




Figure 1: Pipeline for neural predictions for syntax guided program synthesis.

We describe the machine learning pipeline for program synthesis, illustrated in Figure 1 on a toy DSL describing integer list manipulating programs.

The compilation phase constructs a context-free grammar (CFG) from the DSL together with a set of syntactic constraints. The CFG may incorporate important information about the program being generated, such as the n last primitives (encompassing n -gram models) or semantic information (*e.g.* non-zero integer, sorted list).

A prediction model (typically a neural network) takes as inputs a set of I/O and outputs a probabilistic labelling of the CFG, inducing a probabilistic context-free grammar (PCFG). The network is trained so that most likely programs (with respect to the PCFG) are the most likely to be solutions, meaning map the inputs to corresponding outputs.

We refer to Appendix A for an in-depth technical discussion on program representations and on the compilation phase. In this work we focus on the search phase and start with defining a theoretical framework for analysing search algorithms.

The PCFG obtained through the predictions of the neural network defines a probabilistic distribution \mathcal{D} over programs. We make the theoretical assumption that the program we are looking for is actually sampled from \mathcal{D} , and construct algorithms searching through programs which find programs sampled from \mathcal{D} as quickly as possible. Formally, the goal is to minimise the expected number of programs the algorithm outputs before finding the right program.

We write $A(n)$ for the n^{th} program chosen by the algorithm A ; since A may be a randomised algorithm $A(n)$ is a random variable. The performance $\mathcal{L}(A, \mathcal{D})$ of the algorithm A , which we call its loss, is the expected number of tries it makes before finding x :

$$\mathcal{L}(A, \mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}} [\inf \{n \in \mathbb{N} : A(n) = x\}].$$

An algorithm A^* is ‘loss optimal’ if $\mathcal{L}(A^*, \mathcal{D}) = \inf_A \mathcal{L}(A, \mathcal{D})$. Let us state a simple fact: an algorithm is loss optimal if it generates each program once and in

non-increasing order of probabilities. Depending on \mathcal{D} constructing an efficient loss optimal algorithm may be challenging, pointing to a trade off between quantity and quality: is it worth outputting a lot of possibly unlikely programs quickly, or rather invest more resources into outputting fewer but more likely programs?

An example. To illustrate the definitions let us consider the distribution \mathcal{D} over the natural numbers such that $\mathcal{D}(n) = \frac{1}{2^{n+1}}$; it is generated by the following PCFG:

$$S \rightarrow^{\cdot 5} f(S) \quad ; \quad S \rightarrow^{\cdot 5} x,$$

when identifying n with the program $f^n(x)$. Let us analyse a few algorithms.

- The algorithm A_1 enumerates in a deterministic fashion the natural numbers starting from 0: $A_1(n) = n$. Then $\mathcal{L}(A_1, \mathcal{D}) = \sum_{n \geq 0} \frac{n+1}{2^{n+1}} = 2$. This enumeration algorithm A_1 is loss optimal.
- The algorithm A_2 samples the natural numbers using the distribution \mathcal{D} . For $n \geq 0$, the value of $\mathbb{E}[\inf \{n' : A_2(n') = n\}]$ is 2^{n+1} : this is the expectation of the geometric distribution with parameter $\frac{1}{2^{n+1}}$. Then $\mathcal{L}(A_2, \mathcal{D}) = \sum_{n \geq 0} \frac{2^{n+1}}{2^{n+1}} = \infty$. Hence the naive sampling algorithm using \mathcal{D} has infinite loss.
- The algorithm A_3 samples the natural numbers using a distribution that we call $\sqrt{\mathcal{D}}$ defined¹ by $\sqrt{\mathcal{D}}(n) = \frac{1}{1+\sqrt{2}} \frac{1}{2^{\frac{n+1}{2}}}$.

For $n \geq 0$, the value of $\mathbb{E}[\inf \{n' : A_3(n') = n\}]$ is $(1 + \sqrt{2})2^{\frac{n+1}{2}}$: this is the expectation of the geometric distribution with parameter $\frac{1}{1+\sqrt{2}} \frac{1}{2^{\frac{n+1}{2}}}$.

Then

$$\begin{aligned} \mathcal{L}(A_3, \mathcal{D}) &= \sum_{n \geq 0} \frac{(1+\sqrt{2})2^{\frac{n+1}{2}}}{2^{n+1}} \\ &= (1 + \sqrt{2}) \sum_{n \geq 0} \frac{1}{2^{\frac{n+1}{2}}} \\ &= (1 + \sqrt{2})^2 \approx 5.83. \end{aligned}$$

As we will prove in a more general statement (Theorem 2), the algorithm A_3 is loss optimal among sampling algorithms. Surprisingly it is not much worse than the loss optimal algorithm, yet offers many advantages: it is much easier to implement, and requires no memory at all. Last but not least in the case of PCFG it can be implemented using a new probabilistic labelling of the PCFG inducing \mathcal{D} .

The next two sections are devoted to the two natural classes of algorithms for distribution-based search: *enumeration* and *sampling*. We then study how they can run at scale across parallel compute environments.

3 Enumerative methods and the Heap Search algorithm

A number of enumerative methods have been investigated in previous works Menon et al. (2013); Balog et al. (2017); Feng et al. (2018); Zohar and Wolf (2018). They

¹For the normalisation factor, note that $\sum_{n \geq 0} \frac{1}{2^{\frac{n+1}{2}}} = 1 + \sqrt{2}$.




Figure 2: Illustration of the tree of leftmost derivations.

proceed in a top-down fashion, and can be understood as ways of exploring the tree of leftmost derivations of the grammar as illustrated in Figure 2.

We present a new efficient and loss optimal algorithm called **HEAP SEARCH** and following a bottom-up approach. It uses a data structure based on heaps to efficiently enumerate all programs in non-increasing order of the probabilities in the PCFG.

Let us write \mathcal{D} for the distribution induced by a PCFG. For a program x , we say that x' is the ‘successor of x ’ if it is the most likely program after x , meaning $\mathcal{D}(x) > \mathcal{D}(x')$ and there are no programs x'' such that $\mathcal{D}(x) > \mathcal{D}(x'') > \mathcal{D}(x')$. For a non-terminal T in the grammar, the ‘successor of x from T ’ is the most likely program after x among those generated from T . We define ‘predecessor of x ’ and ‘predecessor of x from T ’ in a similar way.

We create a procedure **Query**(T, x) which for any program x generated from a non-terminal T outputs the successor of x from T . Note that once this is defined, the algorithm performs successive calls to **Query**(S, x) with S the initial non-terminal and x the latest generated program, yielding all programs in non-increasing order of the probabilities.

To explain how **Query**(T, x) works, we first define the data structure. For each non-terminal T we have a hash table SUCC_T which stores the successors of all previously seen programs generated from T , and a heap HEAP_T which contains candidate programs, ordered by non-increasing probability. The key invariant is the following: the successor of T from x has either already been

computed, hence is in SUCC_T , or is the maximum program in HEAP_T . This means that implementing $\text{Query}(T, x)$ is very simple: it checks whether the successor has already been computed and returns it in that case, and if not it pops the heap. The difficulty is in maintaining the invariant; for this we need to add a number of candidate programs to the heaps. They are obtained by substituting one argument from the returned successor, and propagate this recursively to the corresponding non-terminals.

Theorem 1. *The HEAP SEARCH algorithm is loss optimal: it enumerates every program exactly once and in non-increasing order of probabilities.*

We refer to Appendix C for a complete description of the algorithm with a pseudocode, a proof of Theorem 1, and a computational complexity analysis.

HEAP SEARCH is related to A^* from Feng et al. (2018): they are both loss optimal algorithms, meaning that they both enumerate programs in non-increasing order of probabilities. As we will see in the experiments HEAP SEARCH is better in two aspects: it is faster, and it is bottom-up, implying that program evaluations can be computed along with the programs and avoiding evaluating twice the same (sub)program.

4 Sampling methods and the SQRT Sampling algorithm

A *sampling algorithm* takes random samples from a distribution \mathcal{D}' ; what is both a strength and a weakness is that a sampling algorithm is memoryless: a weakness because the algorithm does not remember the previous draws, which means that it may draw them again, but also a strength because it uses very little space and can be very easily implemented.

In the case of sampling algorithms, we identify algorithms with distributions. The following theorem shows a dichotomy: either there exists a loss optimal sampling algorithm among sampling algorithms, and then it is characterised as the ‘square root’ of the distribution, or all sampling algorithms have infinite loss.

Theorem 2. *Let \mathcal{D} a distribution over a set X . If $\sum_{x \in X} \sqrt{\mathcal{D}(x)} < \infty$, the distribution $\sqrt{\mathcal{D}}$ defined by*

$$\sqrt{\mathcal{D}}(x) = \frac{\sqrt{\mathcal{D}(x)}}{\sum_{y \in X} \sqrt{\mathcal{D}(y)}}$$

is loss optimal among all sampling algorithms. If $\sum_{x \in X} \sqrt{\mathcal{D}(x)} = \infty$, for all sampling algorithms \mathcal{D}' we have $\mathcal{L}(\mathcal{D}', \mathcal{D}) = \infty$.

Proof. Let \mathcal{D}' be a distribution. For an element x , the expectation of the number of tries for \mathcal{D}' to draw x is $\frac{1}{\mathcal{D}'(x)}$: this is the expectation of success for the geometric distribution with parameter $\mathcal{D}'(x)$. It follows that

$$\mathcal{L}(\mathcal{D}', \mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}} \left[\frac{1}{\mathcal{D}'(x)} \right] = \sum_{x \in X} \frac{\mathcal{D}(x)}{\mathcal{D}'(x)}.$$

Let us assume that $\sum_{x \in X} \sqrt{\mathcal{D}(x)} < \infty$. Thanks to Cauchy-Schwarz inequality we have:

$$\begin{aligned} \left(\sum_{x \in X} \sqrt{\mathcal{D}(x)}\right)^2 &= \left(\sum_{x \in X} \sqrt{\frac{\mathcal{D}(x)}{\mathcal{D}'(x)}} \sqrt{\mathcal{D}'(x)}\right)^2 \\ &\leq \left(\sum_{x \in X} \frac{\mathcal{D}(x)}{\mathcal{D}'(x)}\right) \cdot \underbrace{\left(\sum_{x \in X} \mathcal{D}'(x)\right)}_{=1} \\ &= \sum_{x \in X} \frac{\mathcal{D}(x)}{\mathcal{D}'(x)}. \end{aligned}$$

We note that $\mathcal{L}(\sqrt{\mathcal{D}}, \mathcal{D}) = \left(\sum_{x \in X} \sqrt{\mathcal{D}(x)}\right)^2$, so the previous inequality reads $\mathcal{L}(\mathcal{D}', \mathcal{D}) \geq \mathcal{L}(\sqrt{\mathcal{D}}, \mathcal{D})$. Thus $\sqrt{\mathcal{D}}$ is loss optimal among sampling algorithms, and if it is not defined, then for any \mathcal{D}' we have $\mathcal{L}(\mathcal{D}', \mathcal{D}) = \infty$. \square

Theorem 2 characterises the loss optimal sampling algorithm, but does not explain how to implement it. The following result answers that question.

Theorem 3. *If \mathcal{D} is defined by a PCFG and $\sqrt{\mathcal{D}}$ is well defined, then we can effectively construct a PCFG defining $\sqrt{\mathcal{D}}$.*

The PCFG for $\sqrt{\mathcal{D}}$ is obtained from the PCFG for \mathcal{D} by taking the square root of each transition probability, and then globally renormalising. Details of this procedure can be found in Appendix D.

5 Parallel implementations

Harnessing parallel compute environments is necessary for scalable, future-proof search algorithms, because combinatorial search bottlenecks on compute, and both the present and likely future of massive compute is a parallel one. Accordingly, we have taken care to design and evaluate extensions of our algorithms which can metabolize these compute resources through multiprocessing.

We introduce a new algorithm called the *grammar splitter*, which partitions a PCFG into a balanced family of k sub-PCFGs. Each of the k threads is assigned a sub-PCFG and simply runs a search algorithm on it. Two key advantages of our approach are that any search algorithm can be used in this very simple parallel architecture, and that the theoretical gain of using k threads is linear in k . The output of the grammar splitter is illustrated in Figure 3: the white PCFG is split into 4 sub-PCFGs.

The two crucial properties of the grammar splitter are:

- the space of programs is partitioned into k subspaces. This implies that the threads do not carry out redundant work and that all programs are generated,
- the k program subspaces are balanced, meaning that their mass probabilities are (approximately) equal. This implies that all threads contribute equally to the search effort.

A split is a collection of partial programs, for instance `map (* 2) HOLE` and `fold + HOLE HOLE`, it induces a PCFG. A set of k incomparable splits yields a




Figure 3: The grammar splitter: a balanced partition with quality $\alpha = \frac{.3}{.25} = 1.2$.

partition of the PCFG. Let us write α for the quality of a partition, defined as the ratio between the maximum and the minimum probability mass of a split. We are looking for a balanced partition, *i.e.* one for which the quality α is close to 1.


Our algorithm finds a balanced partition through a hill climbing process: at each point the algorithm either looks for an improving swap or a refinement. In the first case, the action of an improving swap is to transfer a partial program from one split to another, and its goal is to lower the quality coefficient. In the second case, we consider the partial program with maximal probability in a split and refine it: for example `map (* 2) HOLE` could be replaced by `map (* 2) var0` and `map (* 2) (filter HOLE HOLE)`.

6 Experiments

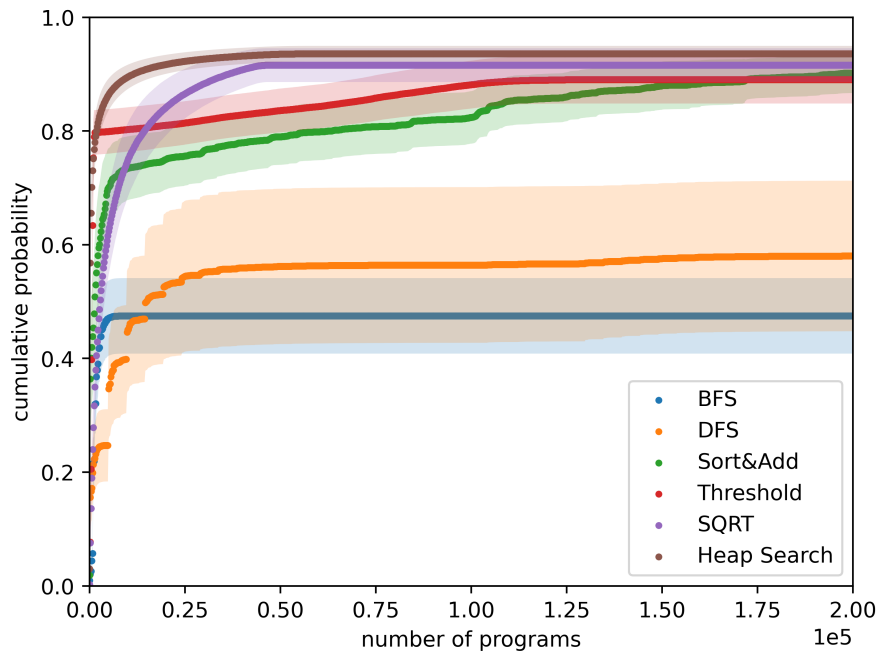
We study a range of search algorithms – both our new ones and prior work – across list processing and string manipulation domains, with the goal of answering the following questions:

- `HEAP SEARCH` and `A*` are both loss optimal enumerative algorithms; beyond these theoretical guarantees, how do the two algorithms compare in practice?
- How effective are our search algorithms for solving complex program synthesis benchmarks using neural guidance?
- How do our algorithms scale with parallel compute?

We use a generic program synthesizer written from scratch in Python (Appendix A), studying random PCFGs (more controlled) and machine-learned



(a) Cumulative probability against time in log-scale



(b) Cumulative probability against number of programs output in log-scale

Figure 4: Comparing all search algorithms on random PCFGs

PCFGs (more naturalistic).

We report results on DSLs from DeepCoder Balog et al. (2017) and DreamCoder Ellis et al. (2021). Both target the classic program synthesis challenge of integer list processing programs, but with different properties. DeepCoder’s DSL is larger and more specialized, with around 40 high-level primitives, and does not use polymorphic types, while DreamCoder’s is smaller and more generic, with basic functional programming primitives such as map, fold, unfold, car, cons, and cdr, etc., for a total of around 20 primitives. Both DSLs are compiled into a CFG with minimal syntactic constraints generating programs of depth 6.

The search algorithms under consideration are:

- THRESHOLD from Menon et al. (2013): iterative-deepening-search, where the threshold that is iteratively deepened is a bound on program description length (i.e. negative log probability),
- SORT AND ADD from Balog et al. (2017): an inner loop of depth-first-search, with an outer loop that sorts productions by probability and runs depth-first-search with the top k productions for increasing values of k ,
- A^* from Feng et al. (2018): best-first-search on the graph of (log probabilities of) tree derivations,
- BEAM SEARCH from Zhang et al. (2018): breadth-first-search with bounded width that is iteratively increased.

As well as our new algorithms: HEAP SEARCH and SQRT SAMPLING. We refer to Appendix B for a description of the algorithms and their implementations.

Our implementation of SQRT SAMPLING uses the Alias method Walker (1977), which is an efficient data structure for sampling from a categorical distribution. We associate to each non-terminal an Alias table, reducing the task of sampling a derivation rule with n choices to sampling uniformly in $[1, n]$ and in $[0, 1]$.

All algorithms have been reimplemented and optimised in the codebase to provide a fair and uniform comparison. We also report on parallel implementations using our grammar splitter.

6.1 Random PCFGs

In this first set of experiments we run all search algorithms on random PCFGs until a timeout, and compare the number of programs they output and the cumulative probability of all programs output.

To obtain random PCFGs from the CFGs we sample a probabilistic labeling with an exponential decrease (this is justified by the fact that machine-learned PCFGs feature exponential decrease in transition probabilities). In this experiment the initialization cost of each algorithm is ignored. The results presented here are averaged over 50 samples of random PCFGs, the solid lines represent the average and a lighter color indicates the standard deviation. Details on the sampling procedure can be found in Appendix F.

Figure 4 shows the results for all algorithms in a non-parallel implementation. On the lhs we see that HEAP SEARCH (almost always) has the highest cumulative probability against both time and number of distinct programs. Note

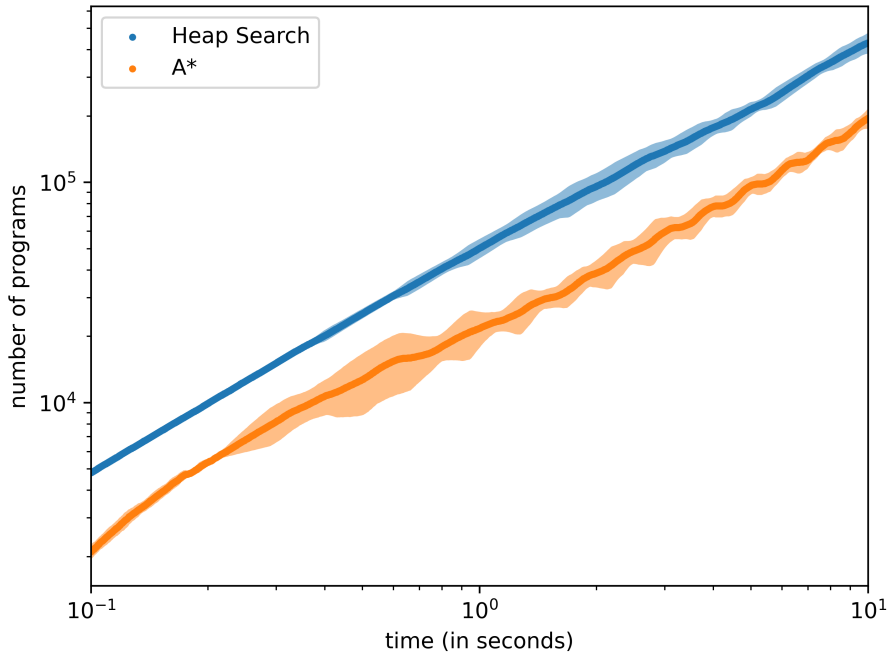


Figure 5: Comparing HEAP SEARCH and A^*

that since A^* and HEAP SEARCH enumerate the same programs in the same order they produce the same curve in the rhs of Figure 4 so we did not include A^* .

To compare A^* and HEAP SEARCH we refer to Figure 5, showing that HEAP SEARCH generates 2.35 times more programs than A^* , consistently over time. The larger variations for A^* are due to the manipulation of a single heap of growing size, requiring frequent memory reallocations. The difference in performance can be explained by the fact that A^* uses a single heap for storing past computations, while HEAP SEARCH distributes this information in a family of connected heaps and hash tables.

We then turn to parallel implementation and perform the same experiments using a variable number of CPUs for HEAP SEARCH and SQRT SAMPLING using the grammar splitter. We do not report on a baseline parallel implementation of SQRT SAMPLING which would simply sample using the same PCFG on multiple CPUs. Indeed this naive approach performs poorly in comparison, since thanks to the grammar splitter two CPUs cannot generate the same program.

The results are shown in Figure 6, where we count programs with repetitions. We see that for SQRT SAMPLING the scale-up is linear, and it is mostly linear for HEAP SEARCH with an acceleration from the 0.2s mark. This acceleration can be explained in two ways: first, each sub-PCFG is shallower since it is split thus it is faster to enumerate program from it, second, once all the successors have been computed HEAP SEARCH is a simple lookup table. At the end of the experiment, SQRT SAMPLING has generated 2.8 times more programs with 6 CPUs than with 2 CPUs, whereas HEAP SEARCH has generated 7.6 times more programs with 6 CPUs than with 2 CPUs.

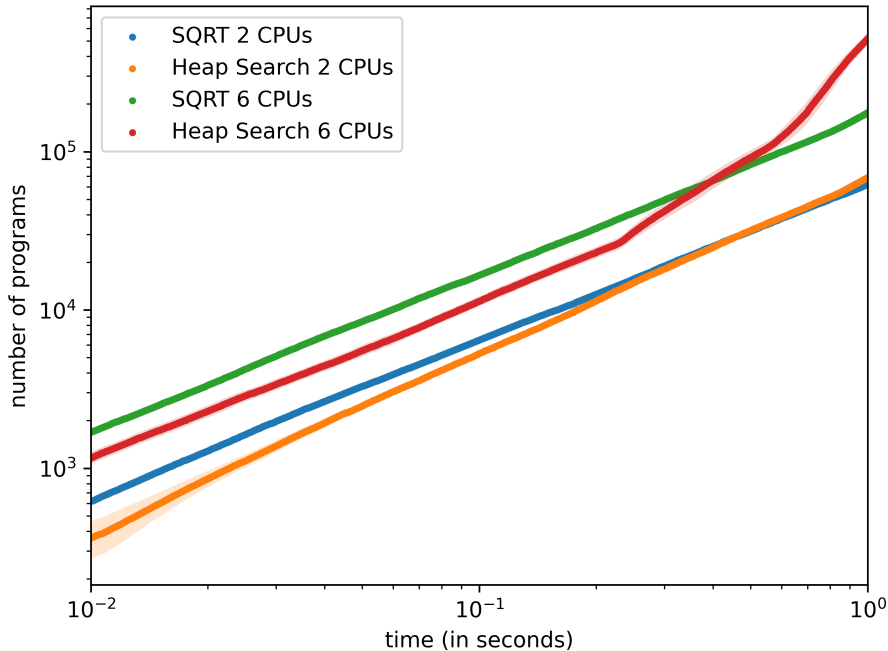


Figure 6: Parallel implementations of HEAP SEARCH and Sqrt SAMPLING using the grammar splitter

This experiment suggests that the grammar splitter enables us to scale our search on multiple CPUs with a linear speed up in the number of CPUs.

6.2 Machine-learned PCFGs

In this second set of experiments we consider the benchmark suites of hand-picked problems and sets of I/O. We extracted 218 problems from DreamCoder’s dataset (Ellis et al., 2021). (The experiments can be easily replicated on DeepCoder’s dataset (Balog et al., 2017) but we do not report on the results here.)

We train a neural network to make predictions from a set of I/O. Our neural network is composed of a one layer GRU Cho et al. (2014) and a 3-layer MLP with sigmoid activation functions, and trained on synthetic data generated from the DSL. The details of the architecture and the training can be found in Appendix F. Our network architecture induces some restrictions, for instance the types of the programs must be `int list -> int list`; we removed tasks that did not fit our restrictions and obtained a filtered dataset of 137 tasks. For each task we run every search algorithm on the PCFG induced by the neural predictions with a timeout of 100s and a maximum of $1M$ programs. Unlike in the previous experiments the initialization costs of algorithms are not ignored.

Figure 7 shows the number of tasks solved within a time budget. HEAP SEARCH solves the largest number of tasks for any time budget, and in particular 97 tasks out of 137 before timeout. The comparison between THRESHOLD and A^* is insightful: A^* solves a bit more tasks than THRESHOLD (85 vs 83) but in twice the time. Sqrt SAMPLING performs just a bit worse than A^* despite

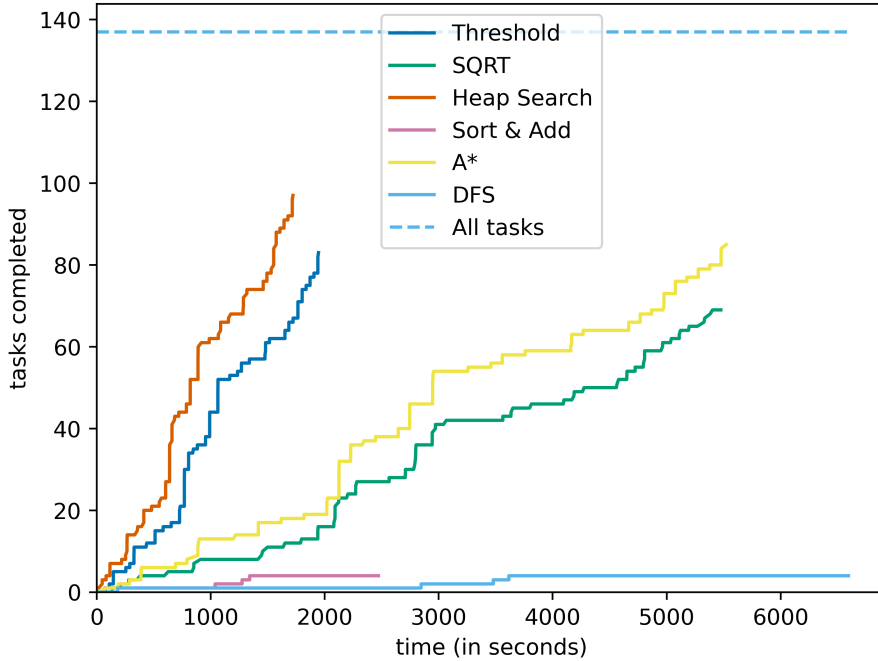


Figure 7: Comparing all search algorithms on the DreamCoder reduced dataset with machine-learned PCFGs

Table 1: Number of programs generated

Algorithm	Number of programs generated
HEAP SEARCH	38735 prog/s
THRESHOLD	25381 prog/s
DFS	20281 prog/s
SQRT SAMPLING	14020 prog/s
A*	6071 prog/s

being a sampling algorithm.

Table 1 shows for each algorithm how many programs were generated per second. Recall that HEAP SEARCH and A* generate the same programs in the same order. Overall in these experiments HEAP SEARCH is 6 times faster than A*.

Since HEAP SEARCH follows a bottom-up approach we save on program evaluations in two ways: partial programs are evaluated along the search, and the results are cached. On the other hand A* is a top-down enumeration method so every new program has to be evaluated from scratch.

It is interesting to compare the rates of SQRT SAMPLING and A*: although SQRT SAMPLING generates over two times more programs, their overall performances are similar. This can be explained in two ways: SQRT SAMPLING may sample the same programs many times, while A* enumerates each program once and starts with the most promising ones according to the predictions.

Finally, we want to see how HEAP SEARCH improves with the quality of the predictions. According to the properties of HEAP SEARCH, we expect that the better the predictions the faster tasks are solved since HEAP SEARCH is loss optimal. In order to show this, we ran HEAP SEARCH on the reduced Dream-Coder dataset with a timeout of 5 minutes and kept the tasks where a solution was found since some tasks cannot be solved with our DSL. Then we trained a neural network on this new reduced set with the solutions found. At different epochs of the training we checked how fast and how many tasks HEAP SEARCH could solve with the predictions of the network with a timeout of 30 seconds and a limit of 1M programs. We plot on Figure 8 the number of tasks solved with respect to total time used by HEAP SEARCH after different number of training epochs with the uniform PCFG as a baseline. We clearly observe the expected outcome, as the number of training epochs grows and the neural networks learn to better predict the solutions, HEAP SEARCH dramatically decreases the time required to solve the tasks. While this may be true for every algorithm, loss optimal algorithms like HEAP SEARCH benefit the most from it.




Figure 8: Evolution of time and tasks solved by HEAP SEARCH at different epochs of the training of a machine-learning PCFG predictor on a reduced DreamCoder dataset

7 Discussion

7.1 Related work

The idea of guiding program search via probabilities is an old one Solomonoff (1989) but which recently has fast become a standard practice in the AI and program synthesis communities. To the best of our knowledge, practical program synthesizers that learn to use probabilistic predictions originated in (Menon et al., 2013), which was first extended with deep learning in (Balog et al., 2017), and such methods are now winning program synthesis competitions Lee et al. (2018). To first approximation, such recent progress has drawn on advances in neural network architecture: e.g., early learned FlashFill-like systems Parisotto et al. (2017) benefited from sophisticated attention mechanisms Devlin et al. (2017), and procedural planning programs Bunel et al. (2018) benefited from feeding execution traces into the neural net Chen et al. (2018). While prior works have explored novel test-time strategies, from Sequential Monte Carlo Ellis et al. (2019) to ensembling methods Chen et al. (2018), here we sought a systematic empirical/theoretical study of two different families of inference strategies, which we intend to mesh well with the larger body of work on neural program synthesis. While the algorithms introduced here do not straightforwardly extend to neural autoregressive models (e.g. RobustFill Devlin et al. (2017)), methods such as SQRT SAMPLING in principle apply to this setting too. We hope that our work here spurs the development of the right tricks to get the theoretical benefits of SQRT SAMPLING for these more flexible model classes, just as DeepCoder paved the way for RobustFill. Many extensions: Lee et al. (2018); Zhang et al. (2018); Polosukhin and Skidanov (2018); Kalyan et al. (2018); Zohar and Wolf (2018).

Shi et al. (2020b) introduced Unique Randomizer in order to sample without replacement: it is a general technique effectively turning a sampling method into an enumerative one by updating the probabilistic weights along the search. It is further improved through batching via Stochastic Beam Search Kool et al. (2019). It is possible to combine the SQRT SAMPLING algorithm with the Unique Randomizer and Stochastic Beam Search. Our experiments did not yield interesting results in that direction, possibly because of memory allocation issues. We leave for future work to optimise this approach.

7.2 Contributions and Outlook

Learning to synthesize programs is a canonical neural-symbolic learning problem: training high capacity statistical models to guide the construction of richly structured combinatorial objects, such as programs. Yet while the neural side of this problem has received much deserved attention, the symbolic component is sometimes taken for granted—after all, symbolic program synthesis has received decades of attention. But the entrance of powerful neural networks for synthesizing programs forces us to reconsider how we deploy symbolic methods for program synthesis. We have considered both systematic and stochastic methods, from both theoretical angles (obtaining guarantees) and also engineering perspectives (such as how to parallelize our new techniques). We hope this work helps contribute to thinking through the symbolic search back-end in this more modern context.

References

- R. Alur, R. Bodík, 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 *Formal Methods in Computer-Aided Design, FMCAD*, 2013. URL <http://ieeexplore.ieee.org/document/6679385/>.
- R. Alur, R. Singh, D. Fisman, and A. Solar-Lezama. Search-based program synthesis. *Communications of the ACM*, 61(12), 2018. URL <https://doi.org/10.1145/3208071>.
- M. Balog, A. L. Gaunt, M. Brockschmidt, S. Nowozin, and D. Tarlow. Deepcoder: Learning to write programs. In *International Conference on Learning Representations, ICLR*, 2017. URL <https://openreview.net/forum?id=ByldLrqlx>.
- R. Bunel, M. J. Hausknecht, J. Devlin, R. Singh, and P. Kohli. Leveraging grammar and reinforcement learning for neural program synthesis. In *International Conference on Learning Representations, ICLR*, 2018. URL <https://openreview.net/forum?id=H1Xw62kRZ>.
- X. Chen, C. Liu, and D. Song. Execution-guided neural program synthesis. In *International Conference on Learning Representations, ICLR*, 2018.
- Z. Chi. Statistical properties of probabilistic context-free grammars. *Computational Linguistics*, 25(1), 1999. URL <https://www.aclweb.org/anthology/J99-1004>.
- 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. In *Conference on Empirical Methods in Natural Language Processing, EMNLP*, 2014. URL <https://aclanthology.org/D14-1179>.
- J. Clymo, A. Gascón, B. Paige, N. Fijalkow, and H. Manukian. Data generation for neural programming by example. In *International Conference on Artificial Intelligence and Statistics, AI&STATS*, 2020. URL <http://proceedings.mlr.press/v108/clymo20a.html>.
- J. Devlin, J. Uesato, S. Bhupatiraju, R. Singh, A. Mohamed, and P. Kohli. Robustfill: Neural program learning under noisy I/O. In *International Conference on Machine Learning, ICML*, volume 70 of *Proceedings of Machine Learning Research*, 2017. URL <http://proceedings.mlr.press/v70/devlin17a.html>.
- K. Ellis, M. Nye, Y. Pu, F. Sosa, J. Tenenbaum, and A. Solar-Lezama. Write, execute, assess: Program synthesis with a REPL. In *Neural Information Processing Systems, NeurIPS*, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/50d2d2262762648589b1943078712aa6-Abstract.html>.
- K. Ellis, C. Wong, M. I. Nye, M. Sablé-Meyer, L. Morales, L. B. Hewitt, L. Cary, A. Solar-Lezama, and J. B. Tenenbaum. Dreamcoder: bootstrapping inductive program synthesis with wake-sleep library learning. In *International Conference on Programming Language Design and Implementation, PLDI*, 2021. URL <https://doi.org/10.1145/3453483.3454080>.

- Y. Feng, R. Martins, O. Bastani, and I. Dillig. Program synthesis using conflict-driven learning. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*, 2018. URL <https://doi.org/10.1145/3192366.3192382>.
- S. Gulwani, O. Polozov, and R. Singh. Program synthesis. *Foundations and Trends in Programming Languages*, 4(1-2), 2017. URL <https://doi.org/10.1561/25000000010>.
- A. Kalyan, A. Mohta, O. Polozov, D. Batra, P. Jain, and S. Gulwani. Neural-guided deductive search for real-time program synthesis from examples. In *International Conference on Learning Representations, ICLR*, 2018. URL <https://openreview.net/forum?id=rywDjg-RW>.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations, ICLR*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- W. Kool, H. Van Hoof, and M. Welling. Stochastic beams and where to find them: The gumbel-top-k trick for sampling sequences without replacement. In *International Conference on Machine Learning, ICML*, 2019. URL <https://proceedings.mlr.press/v97/kool19a.html>.
- W. Lee, K. Heo, R. Alur, and M. Naik. Accelerating search-based program synthesis using learned probabilistic models. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*, 2018. URL <https://doi.org/10.1145/3211992>.
- A. K. Menon, O. Tamuz, S. Gulwani, B. W. Lampson, and A. Kalai. A machine learning framework for programming by example. In *International Conference on Machine Learning, ICML*, 2013. URL <http://proceedings.mlr.press/v28/menon13.html>.
- E. Parisotto, A. Mohamed, R. Singh, L. Li, D. Zhou, and P. Kohli. Neuro-symbolic program synthesis. In *International Conference on Learning Representations, ICLR*, 2017. URL <https://openreview.net/forum?id=rJ0JwFcex>.
- I. Polosukhin and A. Skidanov. Neural program search: Solving programming tasks from description and examples. In *International Conference on Learning Representations, ICLR*, 2018. URL <https://openreview.net/forum?id=BJTtWDyPM>.
- K. Shi, D. Bieber, and R. Singh. TF-coder: Program synthesis for tensor manipulations. In *Workshop on Computer-Assisted Programming, CAP*, 2020a. URL <https://openreview.net/forum?id=nJ5Ij53umw2>.
- K. Shi, D. Bieber, and C. Sutton. Incremental sampling without replacement for sequence models. In *International Conference on Machine Learning, ICML*, 2020b. URL <https://proceedings.mlr.press/v119/shi20a.html>.
- R. Solomonoff. A system for incremental learning based on algorithmic probability. In *Israeli Conference on Artificial Intelligence, Computer Vision and Pattern Recognition*, 1989. URL <https://theworld.com/~rjs/publications/IncLrn89.pdf>.

- A. J. Walker. An efficient method for generating discrete random variables with general distributions. *ACM Transactions on Mathematical Software*, 3(3), 1977. URL <https://doi.org/10.1145/355744.355749>.
- L. Zhang, G. Rosenblatt, E. Fetaya, R. Liao, W. E. Byrd, R. Urtasun, and R. S. Zemel. Leveraging constraint logic programming for neural guided program synthesis. In *International Conference on Learning Representations, ICLR*, 2018. URL <https://openreview.net/forum?id=HJIHtIJvz>.
- A. Zohar and L. Wolf. Automatic program synthesis of long programs with a learned garbage collector. In *Neural Information Processing Systems, NeurIPS*, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/390e982518a50e280d8e2b535462ec1f-Abstract.html>.

A Machine-learned program synthesis

A.1 Program representation as trees

Programs are internally represented as trees as illustrated on the lhs of Figure 9. The program described by this tree is `(take (head (rev l1)) (rev l2))`. We enforce a simple type system for programs consisting of basic types such as `int`, `bool`, `string`, `float` and two constructors: `list` and `arrow`. Each primitive in the DSL comes with a possibly polymorphic type, for instance `take` has the type `take: int -> t0 list -> t0 list`, where `t0` can be instantiated to any type. We only construct correctly typed programs using classical techniques from programming languages such as type inference and de Bruijn’s indices for variables.




Figure 9: Tree and straight line representation of the same program.

A.2 Compilation phase: from domain specific language to grammar

A domain specific language (DSL) is a set of primitives together with their type and semantics. The compilation phase turns a DSL and a set of syntactic constraints into a CFG which generates the set of all correctly typed programs by following typing rules. Syntactic constraints can be very diverse, they include maximum program depth, type request for the program, and other invariants to be satisfied. The type request induces the initial non-terminal of the CFG: for instance if the type request is `int -> int list -> int list` then the initial non-terminal carries the information that the program should output an `int list`, and that we can use one variable of type `int` and another one of type `int list`.

Polymorphic types can be accommodated in two different ways: at compilation time or at generation time. DreamCoder’s implementation Ellis et al. (2021) uses polymorphic types in the CFGs, and performs type inference when generating programs. In our tool we instantiate types during the compilation into CFGs, implying that they do not contain polymorphic types. So when adding the primitive `take: int -> t0 list -> t0 list` we create a number of primitives for each type of bounded size, for instance

```
take[int]: int -> int list -> int list.
```

This makes the CFGs larger, but considerably reduces the computation time for generating programs by avoiding type inference algorithms.

The compilation phase consists in expanding typing rules as long as syntactic constraints are satisfied. The CFGs are trimmed to keep their sizes manageable, removing non-productive and non-reachable non-terminals.

B Review of enumerative methods

We describe the different search algorithms constructed in Menon et al. (2013); Balog et al. (2017); Feng et al. (2018); Zhang et al. (2018), which can be understood as ways of exploring the tree of leftmost derivations of the grammar.

BFS and DFS for CFGs. The first algorithms we mention do not make use of the weights in the PCFG, they are enumeration algorithms for CFGs. We consider the tree of all leftmost partial derivations starting from the initial non-terminal; each node of this tree is a partial program. There are two classical algorithms for searching this tree. Both maintain a tree of partial derivations. To resolve ambiguity, for each non-terminal we fix an (arbitrary) order on the derivation rules from this non-terminal. The breadth-first search (BFS) expands the *highest* node containing a non-terminal by applying the next derivation rule to the leftmost non-terminal in this node (next with respect to the order on derivation rules from this non-terminal). The depth-first search (DFS) instead expands the *lowest* node containing a non-terminal. The DFS may not terminate, we need to introduce a stopping criterion; the most natural is the size of the program: we fix a target size n and before applying a derivation rule we check whether the partial program has size at most n . With this stopping criterion the DFS enumerates all programs of size at most n .

The remaining algorithms take into account the probabilities attached to the PCFG.

The sort and add search Balog et al. (2017). The algorithm is based on two ideas: the first idea is called ‘biassing’; we specify the order on the derivation rules for each non-terminal, naturally by non-increasing order of probabilities, and the second idea is an incremental search: we fix a number k , restrict for each non-terminal to the k most likely derivation rules, runs the search on this restricted grammar, and iterates with a larger k in case the program was not found. The choice of the sequence of values for k is both crucial for the performances and very hard to make. In our experiments we use arithmetic progressions.

The maximal probability from a non-terminal. The following subroutine was introduced in Menon et al. (2013) and will be used by all subsequent algorithms. For each non-terminal T , we write m_T for the most likely program generated from T , and p_T for its probability. Computing m_T and p_T is done using dynamic programming: p_T is the maximum of $p \cdot \prod_{i \in [1, k]} p_{T_i}$ over all derivation rules $T \rightarrow^p f(T_1, \dots, T_k)$.

The threshold algorithm Menon et al. (2013). The algorithm fixes a threshold, enumerates all programs with probability beating the threshold, and iterates with a lower threshold. Enumerating all programs beating a fixed

threshold is achieved using a DFS with the following stopping criterion: before expanding a partial program we check whether it may be completed into a program with probability beating the threshold by replacing each non-terminal T by the probability of the most likely program p_T and checking whether that program has high enough probability. In our experiments we use geometric progressions for the sequence of thresholds.

The beam search Zhang et al. (2018). Beam search is a BFS algorithm with bounded memory: the size of the queue of partial program, called the beam, is restricted to a constant, the beam width.

C Heap Search

The pseudocode is given in Algorithm 1.

Algorithm 1 Heap search

```

1: procedure INITIALIZATION( )
2:   for all non-terminal symbols  $T$  do
3:     Create an empty max heap  $\text{HEAP}_T$ 
4:     Create an empty hash table  $\text{SUCC}_T$ 
5:     Create an empty set  $\text{SEEN}_T$ 
6:     for all derivation rules  $T \rightarrow f(T_1, \dots, T_k)$  do
7:       Add  $f(m_1, \dots, m_k)$  to  $\text{HEAP}_T$  with priority  $\mathcal{D}(f(m_1, \dots, m_k))$ 
8:       Add  $f(m_1, \dots, m_k)$  to  $\text{SEEN}_T$ 
9:     end for
10:  end for
11: end procedure
12:
13: procedure QUERY( $T, x$ )
14:  if  $x$  is a key of  $\text{SUCC}_T$  then
15:    Return  $\text{SUCC}_T[x]$ 
16:  else
17:     $x' \leftarrow \text{pop}(\text{HEAP}_T)$  ▷  $x'$  is the successor of  $x$ 
18:     $\text{SUCC}_T[x] \leftarrow x'$  ▷ updating the data structure
19:    (Assumes  $x' = f(x_1, \dots, x_k)$  is generated by  $T \rightarrow f(T_1, \dots, T_k)$ )
20:    for all  $i \in [1, k]$  do ▷ add all potential successors
21:       $y_i = \text{Query}(T_i, x_i)$ 
22:       $x'_i = f(x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_k)$ 
23:      if  $x'_i$  is not in  $\text{SEEN}_T$  then
24:        Add  $x'_i$  to  $\text{HEAP}_T$  with priority  $\mathcal{D}(x'_i)$ 
25:        Add  $x'_i$  to  $\text{SEEN}_T$ 
26:      end if
27:    end for
28:    Return  $x'$ 
29:  end if
30: end procedure

```

C.1 Correctness proof

The following lemma gives the correctness of the heap search algorithm.

Lemma 1. *For any non-terminal T and program x generated from T , if we have already run $\mathbf{Query}(T,y)$ for any program y preceding x (among those generated from T) then $\mathbf{Query}(T,x)$ outputs the successor of x from T*

Proof. Suppose the lemma is not true and consider the first time where it fails, on say, input (T, x) . At step 14, (T, x) was not queried before and therefore the algorithm branches out to line 16. Saying that the algorithm fails is equivalent to say that line 17 fails, meaning that $x' = \text{pop}(\text{HEAP}_T)$ is not the successor of x from T . Let us denote by y the correct successor. There are two cases which can lead the algorithm the failure:

- **First case:** x' is a program with higher probability than y . In this case, it means that x' was already the output of a previous query (because (T, x) is the first query for which the algorithm fails, so programs with higher probability than y have already been output before). Thus the program x' was pop at least two times from HEAP_T , one time when x' was the correct output of a query, and another time when the program failed: this is not possible since we push programs only once into the heaps thanks to the condition at line 22.
- **Second case:** x' has a smaller probability than the probability of y . In this case, it means that y was not in HEAP_T when the algorithm performed a pop at line 17 (otherwise y would have pop since y has higher probability than x'). Suppose $y = f(x_1, \dots, x_k)$, generated by the derivation rule $T \rightarrow f(T_1, \dots, T_k)$. If all x_i are maximal (meaning that they don't have a predecessor from T_i) then $f(x_1, \dots, x_k)$ is pushed in HEAP_T during the initialization procedure (line 6,7) so the algorithm cannot fail because of this case. Therefore there is at least one x_i , say w.l.o.g x_1 , which has a predecessor x'_1 from T_1 . Consider the program $f(x'_1, x_2, \dots, x_k)$; this program has higher probability than y and therefore has been seen before. Thus $f(\mathbf{Query}(T_1, x'_1), x_2, \dots, x_k)$ was added to HEAP_T because of line 23. To conclude, observe that $f(\mathbf{Query}(T_1, x'_1), x_2, \dots, x_k) = f(x_1, x_2, \dots, x_k)$ so y has previously been added to HEAP_T .

□

C.2 Complexity analysis

Lemma 2. *Fix any non-terminal T . Suppose that we have already generated the first i programs generated from T (meaning that if x is the j -th program for $j < i$, then x is a key of the hash table SUCC_T and $\text{SUCC}_T[x]$ is the $(j + 1)$ -th program generated from T). Then querying the successor of the i -th program has a running time of $O(\log i)$.*

Proof. First observe that a query can call recursively several others queries. However, for any non-terminal symbol T there is at most one query of the form $\mathbf{Query}(T, x)$ which leads the algorithm to branch out to line 16 and thus to possibly rise other queries. Indeed, this case happens only when the successor of x has not been computed yet (otherwise the query stops at line 15); this can

happen for at most one program for any fixed symbol T : the last program from T already seen in any execution of the algorithm. Forgetting about recursive queries, the running time of a query going through line 16 is given by the pop and push operations (line 17 and 23). The number of pops and pushes is at most $m + 1$ where m is the maximal arity of a function in the grammar. Moreover, each such operation costs a running time of $O(\log |\text{HEAP}_T|)$, so the total time for the query is $O(\log |\text{HEAP}_T|)$.

Overall, the total running time is bounded by

$$\sum_T O(\log |\text{HEAP}_T|).$$

To conclude, observe that when we query the successor of the i -th program generated from T , the size of HEAP_T is bounded by $m \cdot i$ since we push at most m programs during any query not already computed before. \square

C.3 Program evaluation

A key advantage of `HEAP SEARCH` is that it is bottom-up: partial programs are composed from the leaves to the root. This implies that partial programs can be evaluated and their evaluations cached. Although memory hungry, this optimisation leads to major gains when taking evaluation into account.

D SQRT Sampling

D.1 An example

The simplest instantiation of this theorem is $X = \{\text{Head}, \text{Tail}\}$ and the Bernoulli distribution \mathcal{D} with parameter p : it draws Head with probability p and Tail with probability $1 - p$.

A sampling algorithm \mathcal{D}' is a Bernoulli distribution with parameter p' , and its loss is

$$\mathcal{L}(\mathcal{D}', \mathcal{D}) = \mathbb{E}_{x \sim \mathcal{D}} \left[\frac{1}{\mathcal{D}'(x)} \right] = \frac{p}{p'} + \frac{1-p}{1-p'}.$$

Minimising for p' yields $p' = \frac{\sqrt{p}}{\sqrt{p} + \sqrt{1-p}}$, inducing the loss optimal sampling algorithm $\sqrt{\mathcal{D}}$.

D.2 Implementation details

The construction follows two steps: first construct a weighted CFG that recognises $\sqrt{\mathcal{D}}$, and then normalise it into a PCFG using Chi (1999). The normalisation requires computing the partition function Z defined by

$$Z(S) = \sum_{P \text{ generated from } S} \mathcal{D}(P).$$

In general the partition function can be computed by solving a system of polynomial equations. This is easier in our case since we restrict ourselves to acyclic PCFGs.

E Parallel implementation

E.1 Description and pseudocode

The pseudocode of the grammar splitter is given in Algorithm 2 using two procedures: split and find improving swap. The split procedure describes at a high level how our grammar splitter works. The find improving swap procedure is here to provide a clear method of finding an improving swap or refinement.

In our experiments, we initialized the splitting as follows: we split the node with highest probability until the total number of nodes is greater than the number of splits required, then assign one node to each split, and the remaining nodes to the last split. We also limited the search of an improving swap or refinement to the most probable split and the least probable split unlike in the find improving swap procedure where all splits are considered. Finally, in all of our experiments $\alpha_{desired} = 1.05$.

F Details on the experiments

F.1 Random PCFG search

To turn the CFGs into PCFGs we sample a probabilistic labelling, meaning a weight for each derivation rule such that the sum over all derivation rules from each non-terminal is one. The distribution depends on a parameter $\alpha \in (0, 1]$: the i^{th} weight is sampled uniformly at random in $[0, \alpha^i]$, and eventually renormalised. The smaller α , the more biased the distribution, implying that the search for programs will be faster since we have better hints on the target program. For $\alpha = 1$ the weights are sampled uniformly, resulting in a very unbiased PCFG, making the search more difficult. In the random PCFGs experiment we use $\alpha = 0.7$.

F.2 Machine-learned PCFG

We only work with tasks of type `int list -> int list`. We remove examples where lists have length greater than $L_{max} = 10$ or where one of the elements of the input or output list are not in $L_{in} = [-30; 30]$.

We say that a task is solved if a program is found which satisfies all examples. The timeout of 100s only takes into account the search time and the evaluation times and not the time to query the neural network for predictions.

The neural network

The neural network takes as input (the encoding of) a list of examples and outputs a probabilistic labelling for the CFG.

Input encoding Each list in the examples is encoded in a naive way by mapping each element of L_{in} to $[0, 60]$. We additionally use two special symbols PAD and NOPAD, leading to a fixed size encoding of a list into a vector of size $2L_{max}$. Hence an example is encoded as a tensor of shape $(n_{inputs}, 2L_{max})$ where n_{inputs} is the number of inputs in the example.

Algorithm 2 Grammar splitter

```
1: procedure SPLIT( $G, nsplits, \alpha_{desired}$ )
2:   Create an initial splitting SPLITS
3:    $\alpha \leftarrow \frac{\max_{sG \in \text{SPLITS}} \text{probability mass}(sG)}{\min_{sG \in \text{SPLITS}} \text{probability mass}(sG)}$ 
4:   while  $\alpha > \alpha_{desired}$  do
5:     if an improving swap exists then
6:       Update SPLITS with the improving swap
7:        $\alpha \leftarrow \frac{\max_{sG \in \text{SPLITS}} \text{probability mass}(sG)}{\min_{sG \in \text{SPLITS}} \text{probability mass}(sG)}$ 
8:     else
9:       Find the partial program with largest probability  $P$ 
10:      Replace  $P$  in its split with its children
11:    end if
12:  end while
13:  Return SPLITS
14: end procedure
15: procedure FIND IMPROVING SWAP( $G, \text{SPLITS}$ )
16:    $\alpha \leftarrow \frac{\max_{sG \in \text{SPLITS}} \text{probability mass}(sG)}{\min_{sG \in \text{SPLITS}} \text{probability mass}(sG)}$ 
17:    $\alpha^* \leftarrow \alpha$  ▷ best improving swap  $\alpha$ 
18:    $s \leftarrow \text{None}$  ▷ best improving swap
19:    $L \leftarrow \text{argmax}_{sG \in \text{SPLITS}} \text{probability mass}(sG)$ 
20:   Sort SPLITS by increasing probability mass
21:   for all  $sG \in \text{SPLITS} \setminus \{L\}$  do
22:     for all  $P' \in L$  do
23:       for all  $P \in G$  do
24:          $\beta \leftarrow \text{Compute new } \alpha \text{ with Swap}(L, sG, P, P')$ 
25:         if  $\beta < \alpha^*$  then
26:            $\alpha^* \leftarrow \beta$ 
27:            $s \leftarrow \text{Swap}(L, sG, P, P')$ 
28:         end if
29:       end for
30:        $\beta \leftarrow \text{Compute new } \alpha \text{ with Gift}(L, sG, P')$ 
31:       if  $\beta < \alpha^*$  then
32:          $\alpha^* \leftarrow \beta$ 
33:          $s \leftarrow \text{Gift}(L, sG, P')$ 
34:       end if
35:     end for
36:   end for
37:    $l \leftarrow \text{argmin}_{sG \in \text{SPLITS}} \text{probability mass}(sG)$ 
38:   Sort SPLITS by decreasing probability mass
39:   for all  $sG \in \text{SPLITS} \setminus \{L, l\}$  do
40:     for all  $P \in sG$  do
41:       for all  $P' \in l$  do
42:          $\beta \leftarrow \text{Compute new } \alpha \text{ with Swap}(l, sG, P, P')$ 
43:         if  $\beta < \alpha^*$  then
44:            $\alpha^* \leftarrow \beta$ 
45:            $s \leftarrow \text{Swap}(l, sG, P, P')$ 
46:         end if
47:       end for
48:        $\beta \leftarrow \text{Compute new } \alpha \text{ with Gift}(sG, l, P)$ 
49:       if  $\beta < \alpha^*$  then
50:          $\alpha^* \leftarrow \beta$ 
51:          $s \leftarrow \text{Gift}(sG, l, P)$  25
52:       end if
53:     end for
54:   end for
55:   Return  $s$ 
56: end procedure
```

Embedding The encoded examples are fed into an embedding consisting of a single layer GRU Cho et al. (2014) outputting a tensor of shape $(s_{GRU} \times 2L_{max})$. In our experiments $s_{GRU} = 10$.

Main layers The embedded examples are fed into an MLP with 3 layers. The first two have output size $s_{MLP} = 64$ and the last layer outputs a tensor of dimension k the number of derivation rules in the CFG C . Assigning these (normalised) weights to the rules of the CFG yields a PCFG.

Training

We optimised end-to-end the neural network with Adam Kingma and Ba (2015) with default parameters and learning rate of $lr = 0.001$. We trained for one epoch with a batch size of 128 on a generated dataset of 10,000 problems. To generate the dataset, programs are sampled from the uniform PCFG and inputs by choosing a length at most L_{max} and elements uniformly at random in L_{in} . If the output of such a generated input does not fall in the lexicon L_{in} then the program is discarded.

We use the binary cross entropy loss between the output of the neural network and the encoding of a solution program: a rule in the CFG has probability 1 if it is used to derive the solution program, and 0 otherwise.