

ProCon: Continuous Enumeration for Just-In-Time Bottom-Up Synthesis

KYLE THOMPSON, ILANA SHAPIRO, and ANIRUDH CANUMALLA, UC San Diego, USA

In order to scale synthesis algorithms to the vast size of the search space without coupling the algorithm with existing training datasets, recent work has combined probabilistic models with just-in-time learning. With just-in-time probabilistically-guided weighted bottom-up-enumerative search, researchers have created an off-the-shelf approach that learns from partial solutions during the synthesis process itself. However, current approaches are based on discrete size-based enumeration, which does not leverage the full power of the probabilistic model since it involves rounding the probabilities to discrete weights. To overcome this limitation, we present *continuous rule-based enumeration*, in which programs are enumerated in order of continuous, nonrounded weights as determined by a given weighting function. We implement our approach in a tool called ProCon, which operates on the syntax-guided synthesis (SyGuS) format and evaluate it on a variety of SyGuS benchmarks. We demonstrate that ProCon's more precise enumeration can solve SyGuS benchmarks while enumerating fewer programs.

CCS Concepts: • **Software and its engineering** → **Domain specific languages; Programming by example;**

Additional Key Words and Phrases: Program Synthesis, Probabilistic Models

ACM Reference Format:

Kyle Thompson, Ilana Shapiro, and Anirudh Canumalla. 2024. ProCon: Continuous Enumeration for Just-In-Time Bottom-Up Synthesis. *J. ACM* 37, 4, Article 111 (March 2024), 11 pages. <https://doi.org/XXXXXXX.XX> XXXXX

1 INTRODUCTION

Syntax-guided synthesis, or SyGuS, is a popular synthesis framework in which the programmer provides a syntactic template for the desired program in addition to the correctness (i.e. semantic) specification. A common instance of SyGuS is a search problem whose input consists of context-free grammar (CFG) that defines the space of possible programs, and the semantic specification comprises of a set of input-output examples. The input-output example paradigm is also called Programming-by-Example, or PBE. From the CFG and input-output examples, the goal of the synthesizer is to find a program generated by the grammar whose behavior is consistent with the semantic specification. The use of inductive semantic specifications such as input-output examples is also known as *inductive program synthesis* [Alur, Bodik, et al. 2013].

One of the primary challenges in synthesis is how to efficiently search for the desired program in the monumentally vast space of possible programs. As the synthesizer considers longer programs, synthesis quickly becomes intractable due to the size of the search space. To improve the scalability of synthesis models, several approaches optimizing synthesis search algorithms have been studied, including EUSOLVER [Alur, Radhakrishna, et al. 2017], EUPHONY [Lee et al. 2018] (which guide

Authors' address: Kyle Thompson, r7thompson@ucsd.edu; Ilana Shapiro, ilshapiro@ucsd.edu; Anirudh Canumalla, acanumalla@ucsd.edu, UC San Diego, 3235 Voigt Dr, La Jolla, CA, USA, 92093.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0004-5411/2024/3-ART111 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

50 top-down search), and TF-CODER [Shi et al. 2022], BUSTLE [Odena et al. 2021], HEAP SEARCH
 51 [Fijalkow et al. 2022], PROBE [Barke et al. 2020], and BEE SEARCH [Ameen and Lelis 2023] (which
 52 guide bottom-up search).

53 Bottom-up enumeration is a dynamic programming technique that maintains a bank of enu-
 54 merated programs and builds new programs by applying production rules to programs from the
 55 bank. A key limitation of most guided bottom-up search methods is how it enumerates programs
 56 during the search. PROBE, for instance, improves its performance by modeling its weight system
 57 after *size-based enumeration*, where the program bank is indexed by AST size, rather than height,
 58 as this has been observed empirically to be more efficient [Barke et al. 2020]. However, in weighted
 59 enumerative search, both height-based and sized-based approaches require discrete values to index
 60 the bank, meaning that transforming these approaches to rank programs by continuous, real-valued
 61 weights is difficult. Thus, PROBE, as well as the similar models in BUSTLE and TF-CODER, rank
 62 programs by rounding real-valued probabilities to discrete costs, i.e. only enumerating programs
 63 *approximately* in the order of decreasing likelihood. This results in a loss of the full power of the
 64 probabilistic model used to guide the search.

65 To overcome this limitation, we propose an alternative approach we call *continuous bottom-up*
 66 *enumeration*. Rather than ranking programs by any discrete metric, we rank *nonterminals* from the
 67 grammar in order of their real-valued continuous weights, allowing us to leverage the full power
 68 of the probabilistic ranking. Specifically, we index the program bank by nonterminals rather than
 69 costs, and ensure that the bank is ordered by continuous weights. We implement our approach
 70 in a tool called PROCON that we build on top of PROBE. Our approach meshes well with PROBE’s
 71 just-in-time-learning paradigm, in which the probabilistic model is learned on the fly. We evaluate
 72 PROCON on 77 string benchmarks and demonstrate that we are able to find the solution with fewer
 73 examples than Probe.

74 **Contributions.** To summarize, this paper makes the following contributions:

- 75 (1) *Continuous bottom-up enumeration*: a new approach to probabilistically-guided bottom-up
 76 enumerative search that enumerates programs in order of real-valued weights (Sec. 3).
- 77 (2) *PROCON*: a prototype implementation of our continuous ranking system with just-in-time-
 78 learning and its evaluation on 77 string benchmarks (Sec. 4).

79 2 RELATED WORK

81 EUSOLVER extends traditional top-down enumerative search approach with a divide-and-conquer
 82 algorithm, enumerating smaller expressions that are correct on subsets of inputs, as well as the
 83 predicates that distinguish these subsets. Then, it combines these expressions and predicates using
 84 a multi-label decision tree learning algorithm to form a conditional expression using Boolean
 85 combinations of the enumerated predicates [Alur, Radhakrishna, et al. 2017]. EUPHONY builds upon
 86 EUSOLVER by extending the grammar with a probabilistic model that determines the likelihood
 87 of each program. EUPHONY achieves superior performance by combining a probabilistic model
 88 with weighted top-down enumerative search (specifically, A^*) to efficiently enumerate programs
 89 in the order of their likelihood. EUSOLVER learns a probabilistic higher order grammar (PHOG)
 90 from known solutions of synthesis problems solved by existing techniques in order to incorporate
 91 a richer context [Lee et al. 2018].

92 More recently, PROBE has drawn from EUPHONY to (1) develop *guided bottom-up-search*, which
 93 uses a PCFG to rank enumerated programs by likelihood via bottom-up, rather than top-down,
 94 enumerative search and (2) use the partial, yet complete, programs encountered during bottom-up
 95 search to enable just-in-time learning of the model. In other words, PROBE learns a PCFG “just-
 96 in-time,” i.e. during synthesis, rather than ahead of time, which eliminates the dependency on
 97 large and often difficult to obtain training datasets that tools like EUPHONY and EUSOLVER require.

99 Programs are enumerated by discrete cost levels determined by the PCFG. The weighted top-down
 100 A* search in EUSOLVER is incompatible with just-in-time-learning, since top-down enumeration
 101 generates incomplete programs that cannot (yet) be evaluated [Barke et al. 2020].

102 TF-CODER, BUSTLE, and HEAP SEARCH take a similar approach: all use a cost function to guide
 103 bottom-up search. The function these systems employ favors programs that are more likely to
 104 lead to a solution. TF-CODER’s cost function requires a manually crafted set of weights for each
 105 operation of the language. During the search, TF-CODER prioritizes combining programs with lower
 106 weights than programs with larger weights, thus biasing the search, where the weight of a program
 107 is defined as the sum of the weights of the production rules used to generate the program [Shi
 108 et al. 2022]. BUSTLE, on the other hand, uses a neural network as a cost function to compute the
 109 probability that a program is part of a solution. The network is a binary classification model that
 110 receives the input-output pairs of the task and the output of a program to each of the input values,
 111 and returns the probability that the program is a subprogram of a solution to the task [Odena et al.
 112 2021]. Finally, unlike TF-CODER, BUSTLE, and PROBE, HEAP SEARCH uses a data structure based
 113 on heaps to efficiently enumerate all programs in non-increasing order of the real-valued PCFG
 114 probabilities. However, it does so with respect to the *evaluation* of the programs, meaning that
 115 each program’s AST must be created to determine its cost. This means that HEAP SEARCH can
 116 generate a very large number of programs that are more expensive than the solution program.
 117 HEAP SEARCH also sacrifices observational equivalence to achieve its best-first ordering.

118 To our knowledge, BEE SEARCH represents the current state-of-the-art in guided bottom-up
 119 enumerative search. Bee Search performs search in a best-first ordering according to novel cost
 120 functions from both the pre-generation (i.e. before AST generation, like PROBE) and post-generation
 121 (i.e. after AST generation, like HEAP SORT) types of cost functions. BEE SEARCH performs best-first
 122 search with respect to the *is* with respect to the *generation* of programs – i.e. Bee Search does
 123 not even create in memory programs that are more expensive than the solution program, thus
 124 circumventing the HEAP SORT problem. BEE SEARCH’s generation-time best-first search is achieved
 125 by searching in an abstract cost-tuple space. One cost-tuple space is defined for each non-terminal
 126 rule, and each cost-tuple state represents a set of programs that are to be generated, such that the
 127 best-first ordering of programs is attained. Bee Search does not sacrifice observational equivalence
 128 checks in its approach. It performs no worse, and usually outperforms, PROBE, TF-CODER, BUSTLE,
 129 and HEAP SEARCH on all its benchmarks, particularly for larger DSLs [Ameen and Lelis 2023].

131 3 IMPLEMENTATION

132 3.1 Problem Description

133 We formulate our problem identically to [Barke et al. 2020]. Our goal is to find programs that
 134 satisfy a given set of input-output examples, \mathcal{E} . Our programs must adhere to a specified grammar,
 135 $\mathcal{G} = (\mathcal{N}, \Sigma, \mathcal{S}, \mathcal{R})$, where \mathcal{N} is a set of nonterminal symbols, Σ is a set of terminal symbols, \mathcal{S} is the
 136 starting nonterminal, and \mathcal{R} is a set of production rules. Any given production rule $R \in \mathcal{R}$ is of
 137 the form $N \rightarrow (t N_1 N_2 \dots N_k)$. We use $N_{\text{rhs}}(R)$ to denote the tuple of nonterminals on the right
 138 hand side of the rule. We use $N_{\text{lhs}}(R)$ to denote the nonterminal on the left hand side of the rule.

140 3.2 Background on PROBE architecture

141 Since PROCON builds on PROBE’s architecture, we first provide an overview of PROBE’s system that
 142 we go on to modify. At a high level, PROBE takes in as input an inductive synthesis problem in
 143 SyGuS format; specifically, a context-free grammar of the desired DSL and a set of input-output
 144 examples. It then alternates between a synthesis phase and a learning phase until a solution is found
 145 or timeout is reached, as seen in Figure 1. This feedback loop enables PROBE’s just-in-time-learning.
 146
 147

In the synthesis phase, PROBE searches over the space of programs in order of increasing discrete costs. During the learning phase, PROBE updates the PCFG using the partial solutions found in the synthesis phase [Barke et al. 2020].

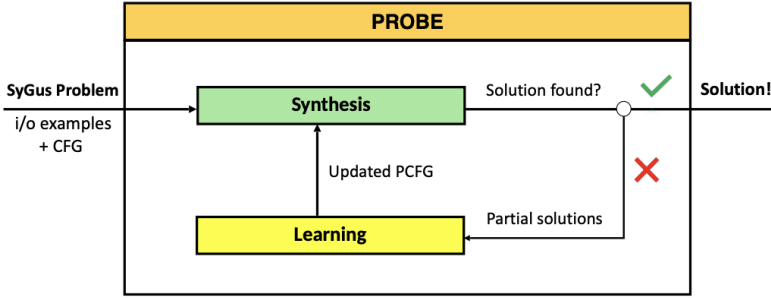


Fig. 1. Overview of PROBE [Barke et al. 2020]

For the synthesis phase, PROBE's algorithm takes a set of input-output examples and enumerates programs in the order of increasing discrete costs according to the PCFG until it finds a program that satisfies the entire specification or reaches a certain cost limit. The algorithm maintains a search state that consists of (1) the current discrete cost level; (2) the program bank storing all enumerated programs indexed by their cost; (3) the evaluation cache storing evaluation results of all programs in the bank in order to check for observational equivalence; and (4) the set of all enumerated partial solutions.

For the learning phase, PROBE uses a simple closed formula to compute new probabilities for each production rule based on the highest proportion of input-output examples that any partial solution derived using this rule satisfies. Since size-based bottom-up enumeration is not amenable to real-valued costs, PROBE enumerates programs *approximately* in the order of decreasing likelihood by converting CFG rule probabilities into discrete costs (the rounded negative logs), which determine the current cost level of that program. The program bank is then indexed by cost, and many programs will end up at the same cost level [Barke et al. 2020].

Every time the PCFG is updated during a learning phase (which happens after the cost level changes), PROBE restarts the bottom-up enumeration from scratch, i.e. it empties the program bank and evaluation cache, and resets the current cost level to zero. Any update to the PCFG renders the program bank outdated, and updating the bank to match the new PCFG requires the amount of computation and/or memory that does not pay off in relation to the simpler approach of restarting the search [Barke et al. 2020].

PROCON only alters the synthesis phase of PROBE's algorithm. It eliminates the need to round the CFG rule probabilities for each program to a discrete cost level, and indeed eliminates the need for cost levels that comprise multiple programs entirely by ranking non-terminals, rather than programs, by their real-valued probabilities.

3.3 Enumeration by Real-Valued Likelihood

In this section, we describe how to perform bottom-up enumeration by real-valued weights. Our algorithm supports any function of the form $W : \mathcal{L}(\mathcal{G}) \rightarrow \mathbb{R}^+$, where $\mathcal{L}(\mathcal{G})$ is the language defined by grammar \mathcal{G} .

Our algorithm uses $W(P)$ to assign a weight to the root node of the input program P . The total weight of a program P is the sum of the weights of all of the nodes in P , denoted $\mathbf{W}(P)$.

For example, consider the following tiny language:

```
S := e
e := arg | (substr e n n)
n := 1 | 0 | (+ n n)
```

Suppose this tiny language had a weight function that partially consisted of

```
W(∅) = 0.1
W(1) = 0.8
W(arg) = 0.5
W((substr arg 0 1)) = 0.3
```

In this case, $W((\text{substr arg } 0 \ 1)) = 1.7$

Our goal is to enumerate programs in order of weights as determined by \mathbf{W} . For this language, that would mean enumerating programs in the following order: \emptyset , arg , 1 , $(\text{substr arg } 0 \ 1)$. Like previous approaches, we maintain a bank of programs that we have already enumerated. We build new programs by combining programs in the bank. Unlike previous approaches that index the bank by integral program costs, we index the bank only by nonterminal. For each nonterminal in the bank, the list of previously enumerated programs is sorted by real-valued weights. This invariant is cheap to maintain as our algorithm enumerates programs by weight, so we can simply append to the end of the bank at the proper nonterminal when we enumerate a new program.

3.3.1 Enumerating Programs for a Particular Rule. We will first describe the algorithm to enumerate programs for a particular production rule $R \in \mathcal{R}$. We enumerate programs from rule R by enumerating subterms to use in the non-terminals of rule R . We enumerate these tuples of subterms in order of the total weight of the tuple. Since the bank of programs is ordered by \mathbf{W} , it is easy to enumerate subterms in order of weight by doing a Dijkstra search over indices into the bank. For sake of example, suppose we want to enumerate subterms to the rule $(\text{substr } e \ n \ n)$. We can achieve this enumeration by enumerating indices into the sequences $\text{bank}[e]$, $\text{bank}[n]$, and $\text{bank}[n]$. It is obvious that the tuple of subterms with the lowest total weight is at indices $(0, 0, 0)$. From here, it is clear that the tuple of subterms with the next lowest weight is at indices $(1, 0, 0)$, $(0, 1, 0)$, or $(0, 0, 1)$. Thus our Dijkstra search forms a tree where nodes are tuples of indices, and each child of a node contains the same tuple of indices with one index incremented. Note that this naive implementation of determining the children of a node could lead to duplicates. This is visualized more clearly in Figure 2. One could take care of duplicates by either maintaining an




Fig. 2. Naive Enumeration with Duplicates

explored set, or by maintaining a bitset of permissions at each node to determine which indices

can be incremented. We opt to maintain a bit set at each node since we hypothesize that it is more efficient. This is visualized in Figure 3, where the red indices are “fixed” and do not have permission to get incremented.




Fig. 3. Enumeration Without Duplicates

Note that this Dijkstra search must happen alongside the growing of the program bank. If we just threw tuples of indices into the frontier, they might be out of bounds of the current program bank. Therefore, we maintain a set of indices, called “abstract candidates” that do not point to anything yet, but will point to something once the bank grows larger. Once an abstract candidate points to concrete subprograms, it can be added to the frontier. The full algorithm for enumerating subprograms to a rule R is given in Appendix A.

3.3.2 Enumerating All Programs. At any particular iteration, to make sure we find the program with the next lowest weight, we must consider the next candidates from each production rule. Specifically, we can keep track of the lowest-weight program so far and enumerate new programs from a production rule while the weight of the **children** of the new programs is smaller than the best candidate. Note that we do the comparison with the children because the enumerator for the production rule is in order of child weight, not program weight. Furthermore, we know that the program weight is greater than or equal to the child weight since $W(\cdot)$ is positive. Once we find the next best program, we add it to the program bank. The full details for enumerating programs in order of likelihood is in Appendix B.

3.3.3 Full ProCon Solver. We combine our enumeration by weight with PROBE’s on-the-fly weight updates in Appendix C. Note that these on the fly weight updates assume that W is in the form of a PCFG. We use a program to update the PCFG if the program satisfies a unique subset of input output examples. We found that it is very important to do these updates in *batches*. One could imagine updating the PCFG after a single program satisfies a unique set of solutions. However, this gives too much power to individual programs and can lead to a noisy search. This is roughly analogous changing the cost level in Probe after just one program is added to Probe’s cost-indexed version of the bank, which would lead to Probe’s PCFG update. We thus set a threshold (`updateWeightLevel` in the pseudocode) leading to a variable sized “batch” of programs to accept based on their weights, before updating the PCFG with this entire batch. Importantly, the threshold is real-valued, since the weights are real-valued.

4 EVALUATION

We evaluated PROCON against 77 string benchmarks with a timeout of 60 seconds. We notice that PROCON solves more benchmarks per program enumerated (Figure 4), but it enumerates programs slower than PROBE (Figure 5). We hypothesize that PROCON is slower at enumerating programs due primarily to the enqueueing and dequeuing in Appendix A. These operations can become expensive since the queue of candidates can grow to the tens of thousands.

295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343

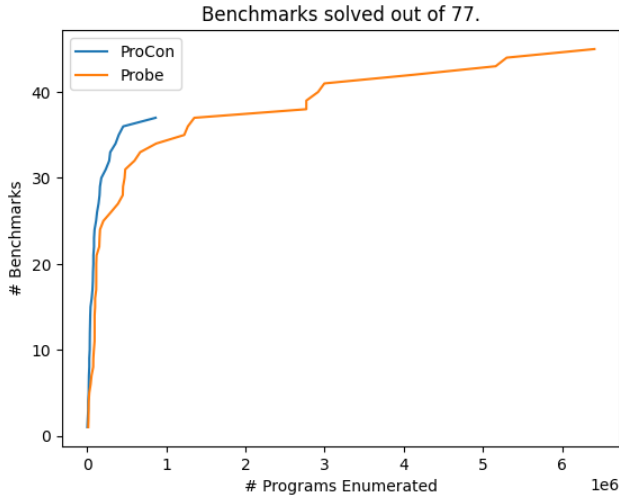


Fig. 4. Evaluation against Probe by number of programs enumerated

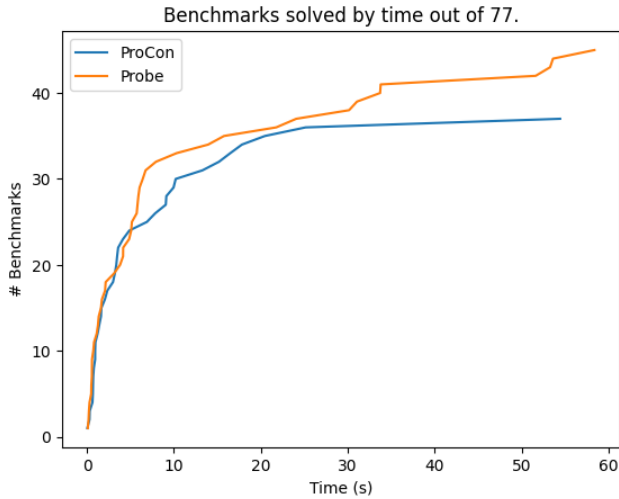


Fig. 5. Evaluation against Probe by wall clock time

5 CONCLUSION AND FUTURE WORK

We have presented an improvement to bottom-up enumerative synthesis with just-in-time learning that leverages the full power of the associated probabilistic model by enabling continuous, real-valued enumeration. We have implemented our algorithm in a tool called PROCON that works with the popular SyGuS input format and uses input-output examples as semantic constraints. We evaluated PROCON on 77 string benchmarks, and found that it was able to arrive at the solution with fewer enumerated programs than PROBE, although there appears to be a computational bottleneck somewhere in our implementation since Probe is still faster.

In future work, we would also like to improve the learning phase of PROCON’s algorithm by exploring other probabilistic models. EUPHONY demonstrated superiority of PHOGs over PCFGs in their approach [Lee et al. 2018], and analogously it seems that PROCON could also be improved by replacing its simple PCFG with a context-aware model. Note that such a model would get context from *subterms* instead of *incomplete programs*. Of course learning a more sophisticated probabilistic model would likely require training data which is currently not required PROCON. One interesting direction could be incorporating a trained probabilistic model with PROBE’s learning on the fly since PROBE’s weights incorporate information from the specification.

We would also like to concretely understand why more-precise enumeration might be better. The learning phase of this algorithm makes it difficult to debug exactly why PROCON comes to a solution before PROBE in a particular example. We would like to run an experiment where we fix the weights of the search to be the weights that lead to the solution. Then, we can run a more direct comparison between continuous and discrete enumeration without worrying about changing weights. Finally, it is important include BEE SEARCH as a baseline in our future evaluations, since they currently represent the state of the art in guided bottom up synthesis.

REFERENCES

- Rajeev Alur, Rastislav Bodik, et al. Oct. 2013. “Syntax-guided synthesis.” *2013 Formal Methods in Computer-Aided Design*, (Oct. 2013). doi: [10.1109/fmcad.2013.6679385](https://doi.org/10.1109/fmcad.2013.6679385).
- Rajeev Alur, Arjun Radhakrishna, and Abhishek Udupa. 2017. “Scaling Enumerative Program Synthesis via Divide and Conquer.” *Tools and Algorithms for the Construction and Analysis of Systems*, 319–336. doi: [10.1007/978-3-662-54577-5_18](https://doi.org/10.1007/978-3-662-54577-5_18).
- Saqib Ameen and Levi H. S. Lelis. 2023. “Program Synthesis with Best-First Bottom-Up Search.” *J. Artif. Intell. Res.*, 77, 1275–1310. doi: [10.1613/JAIR.1.14394](https://doi.org/10.1613/JAIR.1.14394).
- Shraddha Barke, Hila Peleg, and Nadia Polikarpova. Nov. 2020. “Just-in-Time Learning for Bottom-Up Enumerative Synthesis.” *Proceedings of the ACM on Programming Languages*, 4, OOPSLA, (Nov. 2020), 1–29. doi: [10.1145/3428295](https://doi.org/10.1145/3428295).
- Nathanaël Fijalkow, Guillaume Lagarde, Théo Matricon, Kevin Ellis, Pierre Ohlmann, and Akarsh Nayan Potta. 2022. “Scaling Neural Program Synthesis with Distribution-Based Search.” In: *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*. AAAI Press, 6623–6630. doi: [10.1609/AAAI.V36I6.20616](https://doi.org/10.1609/AAAI.V36I6.20616).
- Woosuk Lee, Kihong Heo, Rajeev Alur, and Mayur Naik. June 2018. “Accelerating Search-Based Program Synthesis using Learned Probabilistic Models.” *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*, (June 2018). doi: [10.1145/3192366.3192410](https://doi.org/10.1145/3192366.3192410).
- Augustus Odena, Kensen Shi, David Bieber, Rishabh Singh, Charles Sutton, and Hanjun Dai. 2021. “BUSTLE: Bottom-Up Program Synthesis Through Learning-Guided Exploration.” In: *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net. <https://openreview.net/forum?id=yHeg4PbFhh>.
- Kensen Shi, David Bieber, and Rishabh Singh. 2022. “TF-Coder: Program Synthesis for Tensor Manipulations.” *ACM Trans. Program. Lang. Syst.*, 44, 2, 10:1–10:36. doi: [10.1145/3517034](https://doi.org/10.1145/3517034).

A RULE ITERATOR PSUEDOCODE

Algorithm 1 Rule Iterator

Input: A rule from the grammar: $R \in \mathcal{G}$

Output: Iterator over programs from R in order of the sum of the weights of the subprograms composed by R .

1: **procedure** SETUP

2: argQueues _{$N \in N_{\text{rhs}}(R)$} \leftarrow bank(N)

3: abstractCandidates \leftarrow $\{((0, 0, \dots, 0), (1, 1, \dots, 1))\}$

4: concreteCandidates \leftarrow []

5:

6: **procedure** UPDATEABSTRACTCANDIDATES

7: realizableCandidates \leftarrow $\{c \in \text{abstractCandidates} \mid \text{The indices are in bounds in the bank.}\}$

8: abstractCandidates \leftarrow abstractCandidates \setminus realizableCandidates

9: **for** $c \in \text{realizableCandidates}$ **do**

10: $p \leftarrow \text{REALIZE}(c)$

11: concreteCandidates.enqueue(p)

12:

13: **procedure** NEWABSTRACTCANDIDATES($c \in C$)

14: newAbstractCandidates \leftarrow \emptyset

15: ($_$, $_$, indices, permissions) \leftarrow c

16: **for** $\{i \in [0..|\text{indices}|] \mid \text{permissions}(i)\}$ **do**

17: newIndices(j) \leftarrow if ($j = i$) then (indices(j) + 1) else indices(j)

18: newPermissions(j) \leftarrow $i \leq j$

19: newAbstractCandidates = newAbstractCandidates \cup $\{(\text{newIndices}, \text{newPermissions})\}$

20: **return** newAbstractCandidates

21:

21: **procedure** PEEK

22: UPDATEABSTRACTCANDIDATES()

23: **return** concreteCandidates.peek

24:

25: **procedure** NEXTPROGRAM

26: UPDATEABSTRACTCANDIDATES()

27: nextCandidate \leftarrow concreteCandidates.dequeue()

28: abstractCandidates \leftarrow abstractCandidates \cup NEWABSTRACTCANDIDATES(nextCandidate)

return nextCandidate

B PROGRAM ITERATOR PSUEDOCODE

442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490

Algorithm 2 Program Iterator

Input: A Grammar \mathcal{G}

Output: Iterator over Programs from \mathcal{G} in order of likelihood.

```

1: procedure SETUP
2:   candidateQueue  $\leftarrow$  []
3:   ruleIterators( $R$ )  $\leftarrow$  Rule Iterator for  $R$ 
4:   bank $_{N \in \mathcal{G}}(N) = []$ 
5:
6: procedure NEXTPROGRAM
7:   best  $\leftarrow$  if ( $|\text{candidateQueue}| = 0$ ) then null else candidateQueue.peek
8:   for  $R \in \mathcal{R}$  do
9:      $r \leftarrow$  ruleIterators( $R$ )
10:    while hasNext( $r$ )  $\wedge$  (best = null  $\vee$  r.peek.childWeight < best.weight) do
11:       $p \leftarrow$  r.nextProgram()
12:      if  $p$  is observationally unique then
13:        candidateQueue.enqueue( $p$ )
14:        best  $\leftarrow$  candidateQueue.head
15:   result  $\leftarrow$  candidateQueue.dequeue()
16:   bank(result.nonterminal).append(result)
17:   return result

```

C PROCON SOLVER PSUEDOCODE

Algorithm 3 ProCon Solver

Input: A Grammar \mathcal{G} a set of I/O examples \mathcal{E} , and an update granularity u .

Output: A program from G that satisfies all examples in \mathcal{E} .

```

1: procedure SOLVE
2:   uniquePartialSolutions  $\leftarrow \emptyset$ 
3:   while True do
4:     enumerator  $\leftarrow$  program iterator for  $G$ 
5:     interestingPrograms  $\leftarrow \emptyset$ 
6:     updateWeightLevel  $\leftarrow u$ 
7:     repeat
8:       nextProg  $\leftarrow$  enumerator.nextProgram()
9:       coverage  $\leftarrow$  examples in  $\mathcal{E}$  staisfied by nextProg
10:      if coverage =  $\mathcal{E}$  then
11:        return nextProg
12:      if coverage  $\notin$  uniquePartialSolutions then
13:        uniquePartialSolutions  $\leftarrow$  uniquePartialSolutions  $\cup$  coverage
14:        interestingPrograms  $\leftarrow$  interestingPrograms  $\cup$  nextProg
15:      if  $u \leq$  nextProg.weight  $\wedge$  interestingPrograms =  $\emptyset$  then
16:        updateWeightLevel  $\leftarrow$  updateWeightLevel +  $u$ 
17:      until  $u \leq$  nextProg.weight  $\wedge$  interestingPrograms  $\neq \emptyset$ 
18:      Call UPDATEPROBABILITIES(interestingPrograms)

```

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009