

Semantic Scaffolds for Pseudocode-to-Code Generation

Ruiqi Zhong Mitchell Stern Dan Klein

Computer Science Division

University of California, Berkeley

{ruiqi-zhong, mitchell, klein}@berkeley.edu

Abstract

We propose a method for program generation based on *semantic scaffolds*, lightweight structures representing the high-level semantic and syntactic composition of a program. By first searching over plausible scaffolds then using these as constraints for a beam search over programs, we achieve better coverage of the search space when compared with existing techniques. We apply our hierarchical search method to the SPoC dataset for pseudocode-to-code generation, in which we are given line-level natural language pseudocode annotations and aim to produce a program satisfying execution-based test cases. By using semantic scaffolds during inference, we achieve a 10% absolute improvement in top-100 accuracy over the previous state-of-the-art. Additionally, we require only 11 candidates to reach the top-3000 performance of the previous best approach when tested against unseen problems, demonstrating a substantial improvement in efficiency.

1 Introduction

Systems that can map from natural language descriptions of tasks or programs to executable code have the potential for great societal impact, helping to bridge the gap between non-expert users and basic automation or full-fledged software development. Accordingly, this area of research has garnered significant interest in recent years, with systems being devised for the translation of natural language specifications into database queries (Wang et al., 2018), if-then programs (Chen et al., 2016), game elements (Ling et al., 2016), and more.

While much of the prior work in executable semantic parsing involves short descriptions being mapped into single-line programs, some tasks have recently been proposed that involve multiple natural language utterances on the input side and full programs on the output side, often reaching tens of

Line	Pseudocode	Code
1	in function main	int main() {
2	n is a long integer 0	long n = 0;
3	while n is less than 0	while (n < 'o') {
4
5	close while scope	}

Other wrong candidates

- while (n < 0) { error: use of undeclared identifier 'o'
- while (n < 'o') error: missing '{'

Figure 1: Pseudocode is translated to code for each line and combined to form a valid program. Certain combinations are invalid due to syntactic and semantic constraints.

lines in length and including non-trivial state manipulation. Examples include the Magic the Gathering and Hearthstone datasets (Ling et al., 2016) derived from trading cards and Java or Python classes implementing their behavior in a game engine, the CONCODE dataset (Iyer et al., 2018) consisting of Java documentation strings and method bodies, and the NAPS and SPoC datasets (Zavershynskiy et al., 2018; Kulal et al., 2019) consisting of pseudocode annotations and source code for programming competition problems.

Past approaches to these large-scale language-to-code tasks have typically employed sequence-based models (Ling et al., 2016) that do not account for structure on the output side, or tree-based models (Allamanis et al., 2015; Rabinovich et al., 2017a; Yin and Neubig, 2017; Hayati et al., 2018; Iyer et al., 2019) that incorporate the syntax but not the semantics of the output domain. However, if we want to generate programs that can be executed successfully, the inclusion of both syntactic and semantic constraints is crucial. As shown in Figure 1, while multiple program fragments may be syntactically correct and represent plausible translations of the corresponding pseudocode, not all of them will lead to executable programs.

To address this, we propose a search procedure based on *semantic scaffolds*, lightweight sum-

maries of higher-level program structure that include both syntactic information as well as semantic features such as variable declarations and scope constraints. See Section 3 for a more formal definition. While these do not encode the full spectrum of constraints used in some formal program synthesis tools (Solar-Lezama, 2009; Gulwani et al., 2017), they strike a balance between utility, speed, and ease of use, offering substantial improvements in system performance without a significant increase in complexity.

In this work we focus on the Search-based Pseudocode to Code (SPoC) dataset (Kulal et al., 2019) due to its challenging multiline programs and availability of input-output test suites to evaluate denotation accuracy. The dataset contains line-level pseudocode annotations for 18,356 C++ programs provided by crowdsourcing workers from Amazon Mechanical Turk. As in the approach of Kulal et al. (2019), we first obtain candidate code fragments for each line using an off-the-shelf neural machine translation system. We then aim to find the highest-scoring combination of fragments that results in a valid program. Although finding the optimal program under this setting is NP-hard when variable usage constraints are introduced (see Section A.3), we can approximate it with a hierarchical beam search. Our algorithm first searches for semantic scaffolds for the program, then assembles fragments together conditioned on these scaffolds. This hierarchical approach speeds up search, produces higher quality variations, and leads to substantial improvements in our system’s final accuracy.

We achieve a new state-of-the-art by solving 55.1% of the test cases within 100 attempts. This represents a 10.4% absolute improvement over the previous best (Kulal et al., 2019), and reaches 81% of our model’s oracle performance. When tested against unseen problems (or crowd-workers), our top 11 (or top 52, respectively) candidates have the same performance as their top 3000 candidates, demonstrating marked gains in efficiency.

We complement our results with a discussion of specific cases in which our semantic scaffolds use global program context to resolve ambiguities in the pseudocode. We also conduct a manual error analysis of 200 failures to better characterize the limitations of our method and suggest possible extensions for future work.

Our contributions are summarized as follows:

- We propose the use of semantic scaffolds to

add semantic constraints to models for long-form language-to-code generation tasks.

- We introduce a hierarchical beam search algorithm that incorporates these constraints, resulting in heightened efficiency, better coverage of the search space, and stronger performance when compared with the standard approach.
- We achieve a new state-of-the-art accuracy of 55.1% on the SPoC pseudocode-to-code dataset.

2 Pseudocode-to-Code Task

In this work, we focus on the SPoC dataset introduced by Kulal et al. (2019).

2.1 Data

This dataset consists of C++ solutions to problems from Codeforces, a competitive programming website, along with the input-output test cases used for each problem to evaluate correctness. It contains 18,356 programs in total with 14.7 lines per program on average. Each line is annotated with a natural language pseudocode description given by a crowd worker from Amazon Mechanical Turk. On average, there are 7.86 tokens per line of code and 9.08 tokens per pseudocode annotation. From the full dataset, 1,752 programs with annotations from unseen crowd workers and 1,820 programs for unseen problems are held out for evaluation. More details can be found in Kulal et al. (2019).

2.2 Task

Suppose the target program has L lines. For each line $l \in [L]$, we are given a natural language pseudocode annotation x_l and an indentation level i_l . Our goal is to find a candidate program y based on $(x_1, i_1), \dots, (x_L, i_L)$ that can solve the given problem (i.e. pass all the test cases) using as few submission attempts as possible. The search efficiency of an algorithm is calculated as the fraction of problems it can solve using a budget of B attempts per problem, where an attempt includes both compiling a candidate program and running the test cases.

As in Kulal et al. (2019), for each pseudocode line x_l , we use an off-the-shelf neural machine translation system to obtain a set of C candidate code pieces $Y_l = \{y_{lc} \mid c \in [C]\}$, where candidate code piece y_{lc} has probability p_{lc} . A full candidate

program y is a concatenation of candidate code pieces, one per line, and has score $p(y)$:

$$y = \text{concat}_{l=1}^L y_{lc_l}, \quad p(y) = \prod_{l=1}^L p_{lc_l}. \quad (1)$$

We aim to find valid high-scoring programs in our search procedure.

3 Combination Constraints

Kulal et al. (2019) propose best-first search as a baseline, which enumerates all complete candidate programs in descending order by score. Using a priority queue, this algorithm can efficiently find the exact top B highest scoring candidates in time $O(L \log(BL))$ per candidate.

However, this approach ignores any dependence between different lines. For example, any of the code piece candidates in Figure 1 could potentially be used in a valid program, but if we naively combine certain subsets of candidates together, the resulting program will be invalid due to the use of undeclared variables or mismatching braces. To solve this problem, we propose to enforce certain **syntactic** and **semantic** constraints when combining candidate code pieces.

3.1 Syntactic Constraints

The candidate program should adhere to the grammatical specification of the target language. However, since incorporating the complete set of C++ grammatical constraints would require significant engineering effort, we instead restrict our attention to the set of “primary expressions” consisting of high-level control structures such as `if`, `else`, `for` loops, function declarations, etc. As shown in Figure 2, we parse the candidate code pieces for each line into a list of primary expression symbols. In order for code pieces from consecutive lines to be used together, there must exist a grammatical derivation that combines their respective symbols. The complete list of primary expression can be found in the appendix; see Tables 6 and 7.

Additionally, some production rules are associated with the start or end of a variable scope block. We require that the number of open scope blocks equals the indentation level i_l for each line l .

3.2 Symbol Table Constraints

Each scope block is associated with a symbol table (Aho et al., 1986) keeping track of the variables that have been declared within that scope or any

containing scopes. We extract the variable names used or declared by each code piece (Figure 3) and ensure that (1) undeclared variables are not used, and (2) variables are not redeclared within the same scope. After checking these constraints, any variables declared by a given code piece will be added to the symbol table associated with the current scope.

These symbol table constraints are based on the **semantic** information of code pieces and are fundamentally different from previous AST-based syntactic constraints for code generation (Rabinovich et al., 2017b; Yin and Neubig, 2017). Formally, any context free grammar that specifies the same constraints requires at least exponential description complexity. We provide a proof adapted from Ellul et al. (2005) in Appendix A.2.

3.3 Syntactic and Semantic Scaffolds

We note two properties of the aforementioned constraints. First, we can efficiently compute whether a program prefix can possibly lead to a full program that satisfies the constraints by using an incremental parser (Ghezzi and Mandrioli, 1979) and checking the symbol tables. Secondly, not all information from a code piece is necessary to verify the constraints. Accordingly, when multiple code piece candidates have the same primary expression symbols and variable declarations and usage, swapping between them would not affect the satisfiability of the constraints. For example, changing from `a += 1` to `a -= 1` will not change a compilable program into a non-compilable one, or vice versa. These two properties will help motivate the hierarchical beam search algorithm introduced in the next section.

More formally, we take the *configuration* $\phi(y_{lc})$ of a line y_{lc} to be the minimal set of features required to verify the above constraints. The *prefix scaffold* $S_{y,l} = [\phi(y_{1c_1}), \phi(y_{2c_2}), \dots, \phi(y_{lc_l})]$ of a program y then contains all the information needed to verify the constraints for the first l lines. We can efficiently compute whether $S_{y,l}$ ¹ is a valid prefix scaffold when $l < L$ and whether $S_{y,L}$ is a valid scaffold for a full program when $l = L$.

¹To keep notation uncluttered, we sometimes use ϕ to denote a configuration, we ignore the subscript y of S when we refer to a general scaffold that is not necessarily associated with a specific program, and we ignore the subscript $l = L$ of S when we refer to the scaffold of a full program.

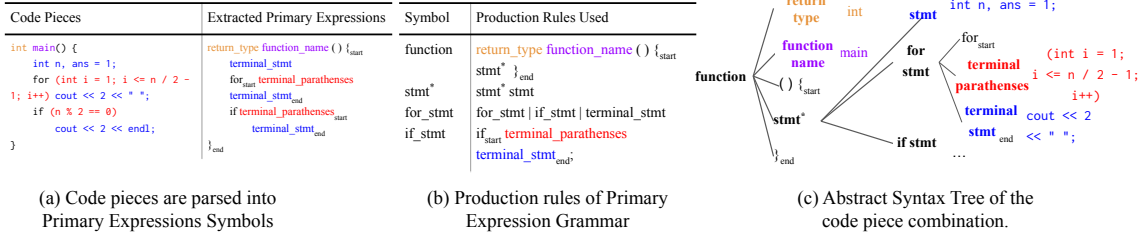


Figure 2: Example primary expression grammar. Subscripts “start/end” refers to starting/ending variable scopes.

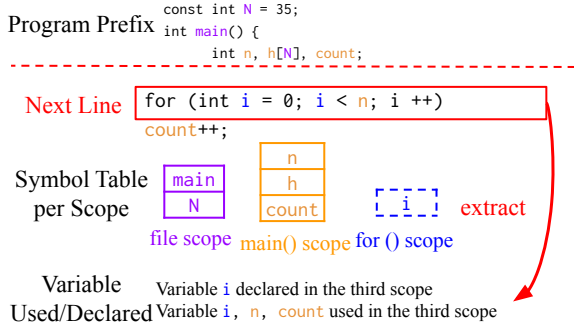


Figure 3: Extracting variables used or declared at each scope for a given code piece to verify the symbol table constraints.

4 Constrained Search

4.1 Beam Search

Our goal is to find the top B highest-scoring candidate programs that satisfy the aforementioned constraints. Unfortunately, finding whether even one solution exists is NP-hard (proof given in Section A.3). One way we can approximate the solution is to use a standard beam search. The beam maintains a list of hypothesis program prefixes along with their respective scores. We extend the beam by adding the candidate code pieces from the next line to each candidate program prefix if they form valid combinations under the constraints, then prune the hypotheses with scores outside of the top W . The algorithm ends after L steps, returning all the valid hypotheses in the final beam.

4.2 Scaffold Search

Although beam search can approximate the top B solutions, the time complexity of beam search grows quadratically with the beam width W . Finding the top B candidates requires that $W \geq B$, and hence each candidate takes $\Omega(BL)$ (amortized) time to generate, which can become intractable if B is on the order of thousands. Even worse, beam search is often biased towards variations at the end of the program due to its greedy decisions, and can

waste its budget on candidates that are unlikely to be the correct solution.

This is in direct contrast to the computationally lighter baseline which generates the exact (unbiased) top candidates independently for each line without constraint. Can we combine the advantages of both algorithms? A key observation is that the assumption of *independent* scoring across different lines allows fast and unbiased full program candidate generation, while an expensive beam search is inevitably needed to deal with the inherent *dependence* between lines.

Therefore, we propose a hierarchical beam search method that first uses beam search with a smaller beam width W to find likely scaffolds, including only the minimum dependency information between lines to satisfy the constraints, then scores candidates independently for each line conditioned on the scaffold. We assign probability $p(\phi_{l\gamma})$ to configuration $\phi_{l\gamma}$ by marginalizing all code piece candidates at line l with configuration $\phi_{l\gamma}$, and assign probability $p(S)$ to scaffold S by multiplying the configuration probabilities from each line:

$$p(\phi_{l\gamma}) = \sum_{\phi(y_{lc})=\phi_{l\gamma}} p_{lc}, \quad p(S) = \prod_{i=1}^L p(S[i]). \quad (2)$$

Using this scoring function, we run a scaffold beam search with size W , then select the top K highest scoring scaffolds $S_1, S_2 \dots S_K$.

Next, to generate program candidates from a given scaffold S , we filter out all code pieces in Y_l that do not have the configuration specified by S ; in other words, the new set of code candidate pieces for each line l is

$$Y_l^S = \{y_{lc} \in Y_l \mid \phi(y_{lc}) = S[l]\}. \quad (3)$$

As a result, conditioned on a fixed scaffold S , code pieces from each line can be chosen independently and the resulting full program will be *guaranteed* to satisfy the aforementioned constraints.

Given K candidate scaffolds, we enumerate the top full program candidate from each scaffold and choose the highest scoring one. This takes time $O(K + L \log(BL))$ per candidate. In practice, we pick relatively small K and the running time has only logarithmic dependence on B .

4.3 Tradeoffs in Early Detection

An alternative view on beam search is that it front loads the computation to reject invalid programs that do not satisfy the constraints earlier in the search process. A brute force alternative is to generate the next highest scoring candidates from the unconstrained baseline and reject invalid ones. This method is guaranteed to produce top-scoring solutions, but it might need arbitrarily many candidates to find a valid one. We need to compare the computational efficiency between these two methods.

The most computationally expensive operation in constraint verification is to verify whether the next line is valid given the program prefix. Therefore, we count how many times this verifier function is called as a proxy to measure computational efficiency. We allow the brute force method to use as large a verifier function call quota as our “active” beam search method: it can validate/reject a program candidate until the quota is used up.

Section 6.4 compares our scaffold search method against this brute force approach. The latter needs thousands of times more computation to attain the same level of performance as the former.

5 Implementation²

Empty Pseudocode Around 26% of the lines in the data set do not have pseudocode annotations. They usually correspond to lines of code that do not have semantically meaningful information, such as “`int main() {`”, “`{`”, “`}`”, etc. Kulal et al. (2019) replaced these empty pseudocode lines with the ground truth code, effectively giving this information away to the search algorithm. We did not use the gold code pieces for these lines, which makes our task more challenging.

Model Training We use OpenNMT (Klein et al., 2017) with its default settings to translate pseudocode into code piece candidates. Our model is a two-layer LSTM seq2seq model with hidden size 512, an attention mechanism (Bahdanau et al., 2014) and copy pointers (Vinyals et al., 2015).

²Our implementation is available at <https://github.com/ruiqi-zhong/SemanticScaffold>.

We estimate the fraction problems solvable given infinite search budget and 100 candidates per line as in Kulal et al. (2019) to obtain an oracle bound on performance. Due to slight difference in hyperparameters and tokenization method, our model has higher ceiling: on the unseen worker (problems) test set, the oracle performance³ is 74.4% (60.5%), compared to 71.4% (55.2%) in previous work. Across all test examples, the oracle performance is 68%.

Parsing Code Pieces Since no off-the-shelf C++ parser extracts the information we need from code pieces, we implement our own primary expression parser to extract high level control information. We rely on the following heuristic assumptions to parse the code pieces generated by the model: (1) a code piece belongs to only one variable scope; (2) the generation of every primary expression terminal symbol lies in one line. Our parser fails on less than 0.01% of the code pieces in the dataset. While selecting the candidates for each line, we immediately reject the ungrammatical pieces we cannot parse. Without deliberate implementation optimization, this parsing operation takes on average 2.6 seconds to process all the top 100 code pieces for a problem – approximately the same wallclock time as 1 compilation attempt.

Search Algorithm Hyperparameters As in Kulal et al. (2019), we consider the top $C = 100$ code pieces for each line. Unless otherwise mentioned, our default beam width W is 50 for scaffold search and we keep the top $K = 20$ scaffolds for the subsequent generation.

6 Search Performance

6.1 Metrics

We evaluate a search algorithm \mathcal{A} by computing the fraction of problem it can solve on the test set given evaluation budget B per problem, which we denote as $f_{\mathcal{A}}(B)$. We plot $f_{\mathcal{A}}$ against B and evaluate it at $B = 1, 10, 100, 1000$ for each algorithm \mathcal{A} to compare performance.

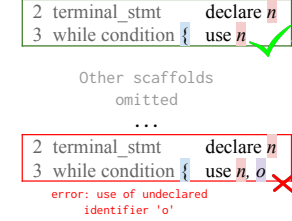
We note that the difference of f values between two algorithms becomes smaller and less informative as B increases. With infinite code piece candidates and budget, a brute force search can

³The oracle performance here is not a universal property of the data, but depends on the model used to generate the code pieces.

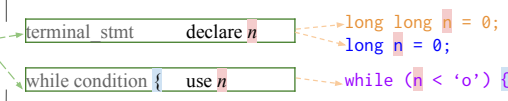
(a) Candidate code pieces and configs ϕ

Line	Pseudocode	Code Piece Candidates	Syntactic Config	SymTable Config
1	in function main	int main() {		
2	n is a long integer 0	long long n = 0; long n = 0;	terminal_stmt ; terminal_stmt ;	declare n declare n
3	while n is less than o	while (n < 'o') { while (n < o) { while (n < 'o')	while condition { while condition { while condition	use n use n, o use n
4	rest of the program omitted ...			

(b) Search over scaffolds



(c) Generate from scaffolds



(d) Combine



Figure 4: (a) Candidate code pieces and their syntactic/Symtable configuration for each line; (b) use beam search to find highest scoring valid scaffolds; (c) given a scaffold, select code pieces that has the same configurations for each line. (d) combine code pieces to form full program.

enumerate all possible programs, find the right solution and f converges to 1. Direct comparison on f values hence becomes meaningless as B increases. To address this deficiency, we define a *lead* metric $l_{\mathcal{A}_1, \mathcal{A}_2}(B)$ equal to the extra budget X needed by algorithm \mathcal{A}_2 to reach the same level of performance as \mathcal{A}_1 given budget B . Formally,

$$l_{\mathcal{A}_1, \mathcal{A}_2}(B) = \inf\{X \mid f_{\mathcal{A}_2}(B + X) \geq f_{\mathcal{A}_1}(B)\}. \quad (4)$$

A visualization can be seen in Figure 5(c).

We report our algorithms' performance on the heldout test set with annotations from unseen crowd workers and with unseen problems separately.

6.2 Comparison of Constraints

We compare four settings:

- **No Constraints:** the best-first search method that scores lines independently.
- **Syntactic Constraints:** the constraints on the primary expression and indentation level as described in section 3.1.
- **Symbol Table Constraints:** both the syntactic constraints and the symbol table constraints described in section 3.2. We abbreviate this as **SymTable**.
- **Backoff:** sometimes hierarchical beam search with the SymTable constraints fails to return

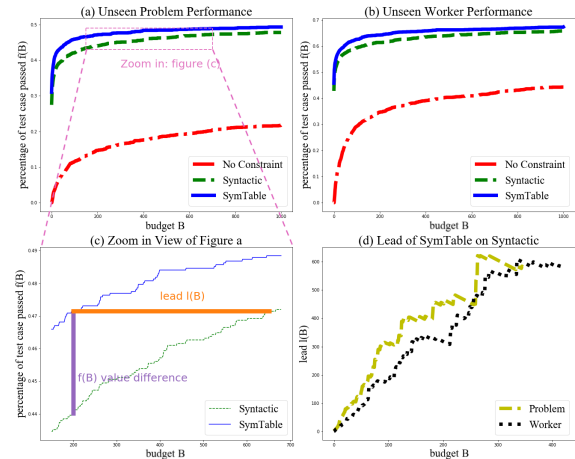


Figure 5: (a), (b) Comparison of f performance under different constraints. (c) a zoom in visualization on the definition of *lead* metrics (d) *lead* of SymTable constraint on Syntactic constraint on different test sets.

any valid scaffold. We back off to just the Syntactic constraints if this happens.

Additionally, we compare with the **Previous** state-of-the-art reported by Kulal et al. (2019).

The results can be seen in Figure 5 and Table 1, where we use the constraint type as a shorthand for the search algorithm under this constraint. Without constraints, the baseline algorithm performs especially poorly because it needs syntactic context to select relevant code pieces for 26% of the lines with empty pseudocode.

SymTable outperforms Syntactic. As shown in

Test Against Unseen Workers				
Hierarchical Search (H), Beam Width $W = 50$				
Constraint	$B=1$	$B=10$	$B=10^2$	$B=10^3$
None	0.0%	8.1 %	29.2 %	44.3%
Previous	30.7%	44.4%	53.7%	58.6%
Syntactic	42.8 %	51.9%	59.3%	65.9%
SymTable	45.8%	55.1%	62.6%	67.3%
Backoff	46.0%	55.3%	62.8%	67.6%
Test Against Unseen Problems				
Constraint	$B=1$	$B=10$	$B=10^2$	$B=10^3$
None	0.0%	3.0%	11.5%	21.8%
Previous	17.8%	28.4%	34.2%	38.3%
Syntactic	27.5 %	35.4%	42.1%	47.8%
SymTable	31.0%	39.2	46.0%	49.3%
Backoff	31.2%	39.4%	46.1%	49.6%

Table 1: Comparison of the fraction of program passed when $B = 10^{0,1,2,3}$ under different constraints; constraint satisfied by hierarchical beam search with the default hyper-parameters mentioned in Section 5. “Previous” refers to the previous state of the art model.

Figure 5(d), the lead of SymTable on Syntactic grows linearly: the more these two algorithms search, the more budget is needed by Syntactic to reach the same level as SymTable. Syntactic needs nearly 600 more budget to have comparable performance with SymTable that uses 400 budget.

We notice that all of our constrained search methods outperform the previous state-of-the-art. Averaged across all test examples, Backoff can solve 55.1% of the problems within 100 budget, which is $\approx 10\%$ higher than the previous work. On unseen workers (problems), the top 11 (top 52) candidates of Backoff solve the same fraction of problems as the top 3000 candidates of the best performing algorithm in Kulal et al. (2019).

6.3 Regular vs. Hierarchical Beam Search

We use regular beam search with beam width $W = 200$ to generate $B = 100$ valid candidate full programs. We did not experiment with $B = 1000$ because beam search with $W \geq B \geq 1000$ is computationally intractable. For hierarchical beam search we experiment with $W = 10, 25, 50$ for scaffold search and keep the top $K = \min(W, 20)$ scaffolds for subsequent searches.

Table 2 compares the performance of hierarchical beam search against regular beam search with different beam sizes under Syntactic and SymTable constraints. We find that if hierarchical beam search is used, even dropping the beam width

Test Against Unseen Workers, Syntactic			
Method, Width	$B=1$	$B=10$	$B=10^2$
H, $W=10$	42.8%	51.7%	59.1%
H, $W=25$	42.8%	51.8%	59.3%
H, $W = 50$	42.8%	51.9%	59.3%
R, $W=200$	42.4%	51.3%	58.2%
Test Against Unseen Workers, SymTable			
Method, Width	$B=1$	$B=10$	$B=10^2$
H, $W=10$	45.4%	54.3%	61.0%
H, $W=25$	45.6%	54.7%	61.9%
H, $W = 50$	45.8%	55.1%	62.6%
R, $W=200$	45.6%	54.9%	61.9%

Table 2: Comparison of different beam size with Syntactic and SymTable constraint when tested against unseen workers. H/R refers to hierarchical/regular beam search and W is the beam width. The same results on unseen problems can be seen in appendix .

from 50 to 10 leads to negligible change in performance. In contrast, even with a large beam width $W = 200$, regular beam search method cannot efficiently search for the solution and leads to a noticeable drop in performance.

We observe a similar trend for SymTable: regular beam search with beam width $W = 200$ underperforms hierarchical search with beam width $W = 25$. However, if we further decrease the hierarchical beam search width from 25 to 10 in this setting, we observe a significant drop in performance, possibly because there are more variable usage variations than syntactic variations.

6.4 Scaffold Search vs. Brute Force Method

We now compare scaffold search to the brute force algorithm as described in section 4.3. We make $B = 50,000$ attempts for the brute force method so that its performance can match at least the top 10 candidates of our constrained approach and make the lead metrics meaningful. To save computation and avoid compiling all 50,000 programs, we early reject every candidate that does not fulfill our constraints.

The lead of our approaches against the brute force algorithm is shown in Figure 6. After being adjusted for the constraint checking quota used, the lead of our approach is tens of thousands ahead of the unconstrained approach. Scaffold search saves lot of computation by inducing a little overhead earlier in the search process.

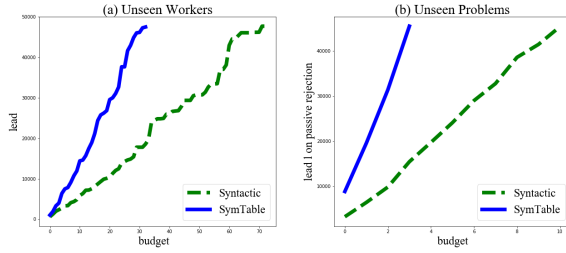


Figure 6: Lead of SymTable and Syntactic constraints on non-constrained approach with equal quota on test set with unseen (a) workers and (b) problems.

7 Analysis

7.1 Program Candidate Variations

Beam search has the problem of producing fewer variations at the beginning of the search. Such a weakness might be tolerable if we only care about the top 1 candidate, but becomes disastrous in a search setting where we want the top B candidates, whose variation is typically spread across the entire program.

We describe the following procedure to formally define this intuition. We first aggregate code piece choices for each line for all the top B programs. As shown in Figure 8(a), we construct a matrix such that each column corresponds to a full program candidate; the number r in the i^{th} row and j^{th} column means that on line i , the j^{th} full program candidate chooses the r^{th} code piece candidate (i.e. $y_{ic_i} = y_{ir}$). Then we can build a prefix tree (Figure 8(b)) by treating each column as a string, where each traversal from the root to a leaf is a complete candidate program y . We define the *representative branch/program* as a traversal from the root to a leaf that always chooses the child that contains the most leaves (with ties being broken randomly). For each of the remaining $B - 1$ programs/traversals, we find the smallest line number where it starts to diverge from the representative branch. Among these $B - 1$ programs, we count the fraction of divergences that take place in the first/second half of the lines. For example, in Figure 8(b), 0% of the divergences occur in the first half.

We compare hierarchical vs. regular beam search under syntactic constraints with different beam widths W : hierarchical $W = 10, 50$ and regular $W = 50, 200$. We group the programs by length L , consider the top $B = 25$ attempted programs for each problem and report the fraction of divergences that occur in the first half of the program length for each group.

Length L	H 10	H 50	R 50	R 200
(0, 10]	45.4%	45.5%	<u>43.6%</u>	45.5%
(10, 20]	63.2%	63.4%	<u>58.2%</u>	63.4%
(20, 30]	63.6%	63.6%	<u>56.8%</u>	63.6%
(30, 40]	67.2%	67.3%	<u>58.2%</u>	67.3%
(40, ∞)	69.4%	68.8%	<u>56.8%</u>	68.8%

Table 3: Fraction of divergence in the first half of the program, grouped by program length L . In the column headers, H/R represents Hierarchical/Regular beam search under **Syntactic** constraint, and the number represents beam width W . The column with the lowest fraction is underlined.

The results can be seen in Table 3. For regular beam search, a moderate beam width $W = 50$ consistently brings fewer variations in the first half of the program, and it needs a larger $W = 200$ to fix this problem. In contrast, a small W for hierarchical beam search produces the same amount of variations in the first half of the program. The same statistics under SymTable constraints can be seen in the appendix (Table 5) and the conclusion holds similarly.

7.2 Rejection by Constraints

In this section we give representative examples on what program candidates are rejected by our syntactic and symbol table constraints.

Syntactic Constraints As mentioned in Section 5, about 26% of the lines do not have pseudocode. They may correspond to “}”, “int main() {”, “{”, “return 0”, “};” or “;”. These lines need contextual information to select valid code pieces and naively combining the top 1 candidate from each line independently will always produce grammatically invalid programs. Syntactic constraints also rule out stylistic ambiguities. For example, when there is only one statement within an `if` statement, the programmer can optionally include a curly brace. However, the pseudocode does not contain such detailed information about style. Both “`if (...) {`” and “`if (...)`” might be valid, but only one of them can be correct given the context of a program. Our syntactic constraints, which contain a curly brace constraint, can help us select the right code piece.

Symbol Table (SymTable) Constraints Pseudocode annotations are sometimes implicit about variable declarations. Given the instruction “set N to 222222”, both code pieces (1) “`int N =`

Reason (percentage)	Pseudocode	Gold Solution	Model Generation
(a) Generation wrong (47.5%)	let value1, value2, val, a, b be integers with val = 0	int value1, value2, val, a, b = 0;	int value1, value2, val = 0, a, b;
(b) Needs type disambiguation (12.5%)	s[occur[i][j] + k] = letter - a + A	s[occur[i][j] + k] = letter - 'a' + 'A';	s[occur[i][j] + k] = 'letter' - a + A;
(c) Needs syntax disambiguation (0.5%)	else if dB is less than dW	} else if (dB < dW) {	else if (dB < dW)
(d) Variable name typos (15.0%)	if lfg = 1	if (flg == 1) {	if (lfg == 1) {
(e) Pseudocode wrong (23.5%)	set ans = 25*length of s	ans += (25 * s.length());	int ans = 25 * s.length();

Figure 7: Categorized error analysis for lines that no generated code piece is functionally equivalent to the gold. The percentage in the parentheses refers to the fraction of this category out of the 200 samples.

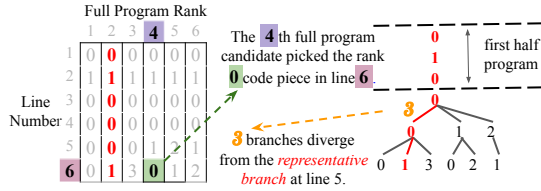


Figure 8: (a) A matrix that represents each candidate’s choices of code pieces for each line. (b) A prefix tree constructed by treating each column as a string; the *representative branch* is the second column and marked with red color.

222222;” and (2) “N = 222222;” are potentially valid. We might disambiguate this case with a SymTable constraint: if the variable is declared before in the same scope, then we know this code piece should not contain a repeated declaration and hence we should choose candidate (2); otherwise we should choose (1) to avoid using undeclared variables. SymTable constraints are also helpful when the pseudocode does not put quotation marks around string/character literals. Consider the instruction “if lucky is A then do the following” with the ground truth code piece “if (lucky == 'A') {”. The model might misunderstand A as a variable name and generate “if (lucky == A) {”. This error can be ruled out by SymTable constraint if variable A is undeclared.

However, SymTable constraints do not preclude all errors related to declarations. Consider the following generation where the last line is wrong:

```
int now = -1, cnt = 0;
for (int i = 0; i < n; ++i) {
    ... // some lines omitted
    // cnt = 1, now = v[i]; // gold
    int cnt = 1, now = v[i]; // pred
}
```

A programmer will usually not declare new variables in the last line of a variable scope. However, technically this is not an invalid statement and the SymTable constraint fails to reject this wrong candi-

date. Extra modelling is needed to take into account programming conventions and common sense.

7.3 Code Piece Error Analysis

So far we have focused on combining independent candidates from each line together to search for the target program. This heavily depends on the underlying model to generate potentially correct code pieces. However, in 32% of the programs at least one “hard” line has no generated code piece that is functionally equivalent to the solution, thus indicating plenty of room for improvement. To help the readers understand the bottleneck for code piece generation and point out important future directions, we randomly sampled 200 “hard” lines and manually analyzed why the generation fails by looking at the top 1 candidate of the model. The error analysis is available on our [GitHub](#).

We group the failures into the following categories, giving a detailed breakdown and examples in Figure 7. (a) The model generation is wrong despite clear pseudocode; this typically happens when the gold code piece is long or highly compositional. (b, c) The pseudocode contains ambiguity; the model generation is reasonable but either needs (b) variable type clarification or (c) syntactic context. This requires incorporating contextual information of the program into the code piece generation process. (d, e) The pseudocode either (d) consists of variable name typos or (e) is completely wrong.

References

Alfred V Aho, Ravi Sethi, and Jeffrey D Ullman. 1986. Compilers, principles, techniques. Addison wesley, 7(8):9.

Miltiadis Allamanis, Daniel Tarlow, Andrew D. Gordon, and Yi Wei. 2015. Bimodal modelling of source code and natural language. In *Proceedings*

- of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15, pages 2123–2132. JMLR.org.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Xinyun Chen, Chang Liu, Richard Shin, Dawn Song, and Mingcheng Chen. 2016. **Latent attention for if-then program synthesis**. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 4581–4589, USA. Curran Associates Inc.
- Keith Ellul, Bryan Krawetz, Jeffrey Shallit, and Mingwei Wang. 2005. Regular expressions: New results and open problems.
- Carlo Ghezzi and Dino Mandrioli. 1979. Incremental parsing. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 1(1):58–70.
- Sumit Gulwani, Oleksandr Polozov, and Rishabh Singh. 2017. **Program synthesis**. *Foundations and Trends® in Programming Languages*, 4(1-2):1–119.
- Shirley Anugrah Hayati, Raphael Olivier, Pravalika Avvaru, Pengcheng Yin, Anthony Tomasic, and Graham Neubig. 2018. Retrieval-based neural code generation. *arXiv preprint arXiv:1808.10025*.
- Srinivasan Iyer, Alvin Cheung, and Luke Zettlemoyer. 2019. Learning programmatic idioms for scalable semantic parsing. *arXiv preprint arXiv:1904.09086*.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2018. Mapping language to code in programmatic context. *arXiv preprint arXiv:1808.09588*.
- G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush. 2017. **OpenNMT: Open-Source Toolkit for Neural Machine Translation**. *ArXiv e-prints*.
- Sumith Kulal, Panupong Pasupat, Kartik Chandra, Mina Lee, Oded Padon, Alex Aiken, and Percy Liang. 2019. Spoc: Search-based pseudocode to code. *arXiv preprint arXiv:1906.04908*.
- Wang Ling, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Andrew Senior, Fumin Wang, and Phil Blunsom. 2016. Latent predictor networks for code generation. *arXiv preprint arXiv:1603.06744*.
- Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017a. Abstract syntax networks for code generation and semantic parsing. *arXiv preprint arXiv:1704.07535*.
- Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017b. **Abstract syntax networks for code generation and semantic parsing**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1139–1149, Vancouver, Canada. Association for Computational Linguistics.
- Armando Solar-Lezama. 2009. **The sketching approach to program synthesis**. In *Proceedings of the 7th Asian Symposium on Programming Languages and Systems, APLAS '09*, pages 4–13, Berlin, Heidelberg. Springer-Verlag.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, pages 2692–2700.
- Chenglong Wang, Po-Sen Huang, Alex Polozov, Marc Brockschmidt, and Rishabh Singh. 2018. **Execution-guided neural program decoding**. *CoRR*, abs/1807.03100.
- Pengcheng Yin and Graham Neubig. 2017. **A syntactic neural model for general-purpose code generation**. In *The 55th Annual Meeting of the Association for Computational Linguistics (ACL)*, Vancouver, Canada.
- Maksym Zavershynskiy, Alex Skidanov, and Illia Polosukhin. 2018. Naps: Natural program synthesis dataset. *arXiv preprint arXiv:1807.03168*.

A Appendices

A.1 Primary Expressions

Table 6 contains the grammar we use for the syntactic constraint and Table 7 defines the generation of terminal symbols.

A.2 CFG Description Size of SymTable

We show that we cannot specify the SymTable constraint in a context free grammar without exponential description complexity w.r.t. the number of variables declared. The intuition is that, since repeated declarations of a variable are not allowed, we need to keep track of all the variables that have been declared every time when verifying whether the next line is valid; however, a CFG, when transformed into a pushdown automata, is only allowed to peek at the top of the stack to decide the state transition. This means the symbol on the top of the stack, the state, or the transition rule need to have full information of about whether each variable has been declared, which contains exponentially many possibilities w.r.t. the number of variables.

Our proof is an adaptation of Ellul et al. (2005), which proves this property for the language that accepts all the permutations of a fixed number of variables. We refer the readers to this paper if more details of the proof are needed. To formalize, we consider a simple grammar of K characters $\{v_1, \dots, v_K\}$, where v_i means, semantically, declaring the variable v_i , and the language L consists of all the possible sequences of declarations that have no repetition.

$$L = \{\text{concat}_{j=1}^k v_{a_j} \mid a_{j_1} \neq a_{j_2} \text{ if } j_1 \neq j_2, k \leq K\} \quad (5)$$

We prove that

Theorem 1 L has at least $\tilde{\Omega}(1.37^K)$ description complexity⁴ as a context free grammar.

Intuitively, it means if we want to use a CFG to specify L , we need the sum of total length of the production rules and number of symbols to be at least exponential.

Proof: Since we can convert any CFG with size B to Chomsky Normal Form (CNF) with size $O(B^2)$, the above statement would be implied if we prove that L needs $\tilde{\Omega}(1.37^{2K}) = \tilde{\Omega}(1.89^K)$ description size in Chomsky Normal Form.

We use Lemma 31 from Ellul et al. (2005):

⁴ $\tilde{\Omega}$ ignores all the poly(K) multiplicative factors.

Lemma 2 Let S be the start symbol of the CFG. Then for all $w \in L$, there exists a symbol A with

$$S \Longrightarrow^* \alpha A \beta \Longrightarrow^* w \quad (6)$$

such that if A yields y in w (i.e. $w = \alpha y \beta$), $\frac{1}{3}|w| \leq |y| \leq \frac{2}{3}|w|$.

In other words, for any member of the language, we can find a symbol in the derivation responsible for between 1/3 and 2/3 of the final yield.

Let P_K be all sequences of permutations of the K variables and thus $P_K \subset L$. Then by Lemma 2, for every permutation $\pi \in P_K$ we can find yield y_π that is yielded by a single symbol such that $\frac{1}{3}K \leq |y_\pi| \leq \frac{2}{3}K$. Now we consider two permutations π_1 and π_2 . If y_{π_1} and y_{π_2} are yielded by the same symbol, then they must have the same length (this is the part where the proof is slightly different from Ellul et al. (2005)): suppose the contrary, w.l.o.g., let $|y_{\pi_1}| > |y_{\pi_2}|$. By the definition of a context free grammar, we can replace the sub-string y_{π_2} in π_2 by y_{π_1} to create a new string y'_{π_2} which is still a member of L . We have $|y'_{\pi_2}| = K - |y_{\pi_2}| + |y_{\pi_1}| > K$ by assumption. However, there are in total K variables; by the pigeonhole principle there must be a variable that is declared twice, and hence $y'_{\pi_2} \notin L$ and we obtain a contradiction.

Then all the assumption needed by Theorem 30 in Ellul et al. (2005) hold and L has description complexity $\tilde{\Omega}(1.89^K)$ in CNF and hence L has description complexity $\tilde{\Omega}(1.89^{K/2}) = \tilde{\Omega}(1.37^K)$. \square

A.3 Hardness of Satisfying SymTable

We show that combining code pieces from each line under the SymTable constraint is NP-Hard in general. We first remind the readers of the set packing problem:

Definition 3 Assume the universe to be \mathcal{V} , and suppose we are given a family of subsets \mathcal{S} from the power set of \mathcal{V} , i.e. $P(\mathcal{V}) = \{S \mid S \subseteq \mathcal{V}\}$ and $\mathcal{S} \subseteq P(\mathcal{V})$. We want to determine whether we can find a packing $\mathcal{K} \subseteq \mathcal{S}$ for which all sets in \mathcal{K} are pairwise disjoint and with size $|\mathcal{K}| \geq L$ for some fixed $L > 0$. This problem is called the set packing problem, and is known to be NP-complete.

Following the notation in section A.2, for each line $l \in [L]$, we construct the $C = |\mathcal{S}|$ code piece candidates y_{lS} for $S \in \mathcal{S}$ as

$$y_{lS} = \text{concat}_{v \in S} v. \quad (7)$$

Test Against Unseen Problems, Syntactic			
Method, Width	$B=1$	$B=10$	$B=10^2$
H, $W=10$	27.4%	35.3%	42.0%
H, $W=25$	27.5%	35.4%	42.1%
H, $W=50$	27.5%	35.4%	42.1%
R, $W=200$	27.1%	34.7%	41.0%
Test Against Unseen Problems, SymTable			
Method, Width	$B=1$	$B=10$	$B=10^2$
H, $W=10$	30.3%	38.1%	43.1%
H, $W=25$	30.9%	39.2%	45.7%
H, $W=50$	31.0%	39.2%	45.9%
R, $W=200$	30.7%	38.9%	45.4%

Table 4: Comparison of different beam size with **Syntactic** and **SymTable** constraint when tested against unseen problems. H/R refers to hierarchical/regular beam search and W is the beam width. This table is structured similarly as 2 .

Length L	H 25	H 50	R 50	R 200
(0, 10]	40.7%	41.5%	39.4%	41.5%
(10, 20]	60.9%	59.8%	54.3%	61.3%
(20, 30]	62.2%	61.3%	54.2%	61.3%
(30, 40]	66.0%	66.1%	56.8%	66.1%
(40, ∞)	69.0%	68.7%	57.9%	68.7%

Table 5: Fraction of divergence in the first half of the program, grouped by program length L . In the column headers, H/R represents Hierarchical/Regular beam search under **SymTable** constraint, and the number represents beam width W .

We easily see that there is a set packing of size L if and only if there is a valid code piece combination under SymTable constraint (declarations need to be disjoint for each line). Hence we finish our reduction proof. \square

A.4 Beam Search on Unseen Problems

Table 4 contains similar information as in Table 2, except that the results are obtained on testing with unseen problems. The exact same conclusion holds: for regular beam search, small beam size hurts performance, but hierarchical beam search can solve this problem.

A.5 Variation under SymTable Constraints

Table 5 contains similar information as Table 3, but for SymTable constraints. The same trend holds: regular beam search with small beam size have fewer variations in the first half of the program.

Symbol	Production Rule
program	stmt program function program
stmt	for_stmt if_stmt while_stmt dowhile_stmt terminal_stmt ;
X^*	$X^* X$ X $\langle \text{EMPTY} \rangle$
function	return_type function_name (args) { _{start} stmt* } _{end} return_type function_name (type*);
args	$\langle \text{EMPTY} \rangle$ arg , args
arg	type arg_name
for_stmt	for _{start} terminal_parentheses terminal_stmt _{end} ; for _{start} terminal_parentheses {stmt*} _{end}
while_stmt	while _{start} terminal_parentheses terminal_stmt _{end} ; while _{start} terminal_parentheses {stmt*} _{end}
dowhile_stmt	do _{start} {stmt*} while terminal_parentheses _{end} ; do _{start} terminal_stmt while terminal_parentheses _{end} ;
if_stmt	single_if_stmt elif_stmt* else_stmt single_if_stmt elif_stmt*
single_if_stmt	if _{start} terminal_parentheses terminal_stmt _{end} ; if _{start} terminal_parentheses {stmt*} _{end}
elif_stmt	elif _{start} terminal_parentheses terminal_stmt _{end} ; elif _{start} terminal_parentheses {stmt*} _{end}
else_stmt	else _{start} terminal_stmt _{end} ; else _{start} {stmt*} _{end}

Table 6: The full primary expression grammar we are using. Each line is a production rule. X is a generic symbol.

Terminal	Implementation
terminal_parentheses	a string that has matching parentheses and starts with parentheses
terminal_stmt	a string that does not contain “;”, “for”, “if”, “else”, “while”, “do”
for, if, else, while, do	reserved key words
function_name, arg_name	function name and function argument name
return_type, type	type in C++

Table 7: The definition of the terminals appearing in Table 6