

Incremental Context-free Grammar Inference in Black Box Settings

Feifei Li^{†‡}, Xiao Chen[§], Xi Xiao^{†‡*}, Xiaoyu Sun[¶], Chuan Chen[‡], Shaohua Wang^{||*}, Jitao Han^{||}

[†]Tsinghua Shenzhen International Graduate School, China

[‡]Key Laboratory of Computing Power Network and Information Security, Ministry of Education, Qilu University of Technology (Shandong Academy of Sciences), China

[§]The University of Newcastle, Australia

[¶]The Australian National University, Australia

^{||}Central University of Finance and Economics, China

ABSTRACT

Black-box context-free grammar inference presents a significant challenge in many practical settings due to limited access to example programs. The state-of-the-art methods, ARVADA and TREEVADA, employ heuristic approaches to generalize grammar rules, initiating from flat parse trees and exploring diverse generalization sequences. We have observed that these approaches suffer from low quality and readability, primarily because they process entire example strings, adding to the complexity and substantially slowing down computations. To overcome these limitations, we propose a novel method that segments example strings into smaller units and incrementally infers the grammar. Our approach, named KEDAVRA, has demonstrated superior grammar quality (enhanced precision and recall), faster runtime, and improved readability through empirical comparison.

ACM Reference Format:

Feifei Li^{†‡}, Xiao Chen[§], Xi Xiao^{†‡*}, Xiaoyu Sun[¶], Chuan Chen[‡], Shaohua Wang^{||*}, Jitao Han^{||}. 2024. Incremental Context-free Grammar Inference in Black Box Settings. In *Proceedings of 39th IEEE/ACM International Conference on Automated Software Engineering (ASE '24)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Inferring context-free grammars (CFG) [9] from example strings [26–28] is fundamental to various software engineering tasks, including code understanding [25], reverse engineering [21], protocol specification [8, 23], detecting and refactoring code smells [15, 22], and generation of randomized test inputs [1, 4, 10, 11, 24, 31, 33]. Traditional methods for grammar inference often rely on grey-box or white-box approaches, where internal structures of parsers are

accessible [7, 13, 19]. However, in many real-world scenarios, only black-box parsers are available, making context-free grammar inference particularly challenging.

In a black-box setting, practitioners are typically provided with a set of example strings that adhere to an underlying grammar. Accompanying these examples is a closed-source parser, often referred to as an oracle, which has the capability to assess and verify the validity of input strings against the established grammar. The primary challenge in this environment is to infer the complete grammar based solely on the limited information provided by the example strings and the binary feedback from the oracle. This task involves constructing a robust and comprehensive representation of the grammar that can accurately predict the oracle's validation outcomes for any given input, essentially replicating the oracle's rule set without access to its internal logic or structure.

Recent advancements in black-box CFG inference include ARVADA [17] and TREEVADA [3]. ARVADA approaches the task by inferring grammar directly from the entirety of input example strings and by exploring various generalization sequences. This method, however, introduces several significant drawbacks: notably, low accuracy, due to difficulties in detecting over-generalization within complex grammars; reduced processing speed; and diminished grammar readability, resulting from complex grammar structures. On the other hand, TREEVADA [3] has made improvements over ARVADA [17] by implementing common language concept nesting rules to pre-structure input example strings, aiming to enhance the structural predictability and efficiency of the inference process. Despite these improvements, TREEVADA [3] still faces the core challenges of low accuracy, slow processing speeds, and limited readability that are inherent in the methods used by its predecessor.

We propose KEDAVRA to address the limitations identified in previous works. Firstly, KEDAVRA incorporates a data decomposition step that simplifies the processing of complex data into more simple units. This enables incremental construction of the grammar, ensuring that our results are robust and less susceptible to variations in the dataset. Such decomposition not only facilitates more stable grammar inference but also enhances processing speed, particularly with complex datasets. Secondly, KEDAVRA employs an incremental approach to grammar inference, systematically constructing the grammar from the simplest to the most complex structures. This method significantly improves the readability of the resulting grammar and simplifies the identification of potential over-generalizations during the inference process. Consequently, the accuracy of the grammar is substantially improved.

*Corresponding authors: Xi Xiao and Shaohua Wang.

Emails: Feifei Li (lff23@mails.tsinghua.edu.cn), Xiao Chen (xiao.chen@newcastle.edu.au), Xi Xiao (xiaox@sz.tsinghua.edu.cn), Xiaoyu Sun (xiaoyu.sun1@anu.edu.au), Chuan Chen (chenchuan2019@qlu.edu.cn), Shaohua Wang (davidshwang@ieee.org), Jitao Han (hanjitao1@gmail.com).

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, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ASE '24, Oct 27– Nov 1, 2024, Sacramento, California, United States

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

ACM ISBN 978-1-4503-XXXX-X/18/06

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

In summary, this paper presents several significant contributions:

- We introduce KEDAVRA, a novel approach for CFG inference in black-box environments. The inferred grammars demonstrate robust stability and exhibit minimal sensitivity to variations in the example strings, ensuring reliable performance across diverse datasets.
- Through rigorous empirical evaluation, we compare the performance of KEDAVRA against the state-of-the-art approaches, ARVADA and TREEVADA. Our results indicate that KEDAVRA significantly outperforms these benchmarks in terms of grammar quality, readability, consistency, and computational efficiency.
- We proposed a more reliable sampling algorithm that can be used to evaluate other grammar inference tools.
- We have made the source code of KEDAVRA and all associated experimental results publicly available¹

2 BACKGROUND

2.1 Context-free grammar

Context-Free Grammar (CFG)[9] is a mathematical model used to describe the grammar structure of formal languages. Formally, a context-free grammar G can be represented as a quadruple (N, Σ, P, S) , where N is a finite set of non-terminal symbols, typically denoted by uppercase letters; Σ is a finite set of terminal symbols, which constitute the actual strings, or the alphabet of the language, typically denoted by lowercase letters or specific characters; P is a finite set of production rules, each rule taking the form $A \rightarrow \alpha$, where A is a non-terminal symbol and α is a string composed of terminals and/or non-terminals; and S is the start symbol, a special non-terminal symbol indicating the beginning of a sentence or the start of the grammar.

Consider a simple context-free grammar defined as follows: the non-terminal symbol set $N = \{S, A\}$, the terminal symbol set $\Sigma = \{a, b\}$, and the production rules $S \rightarrow aA$, $A \rightarrow Sb$, $A \rightarrow ab$, with the start symbol S . Through these rules, strings such as "aab", "aaabb", "aaaabbb", etc., can be generated. The generation process begins with the start symbol S , then applies the rules iteratively until a string of terminal symbols is formed. For instance, starting with S , we can apply the rule $S \rightarrow aA$ to obtain aA , then apply $A \rightarrow Sb$ to get aSb , further apply $S \rightarrow aA$ to yield $aaAb$, and finally apply $A \rightarrow ab$ to generate $aaabb$.

The importance of context-free grammars lies in their sufficiently powerful expressiveness to represent the grammar of most programming languages; in fact, almost all programming languages are defined using context-free grammars. On the other hand, context-free grammars are simple enough that we can construct efficient parsing algorithms to verify whether a given string is generated by a specific context-free grammar. In real-world applications, it could be used to validate the grammar of some structured data formats, such as JSON[30] and XML[20]. In addition, compilers or interpreters of programming languages such as C++[18], Java[12] and JavaScript[37] can use it to check whether the code conforms to the grammar when compiling.




Figure 1: A simple example of ARVADA workflow

2.2 Black-box Grammar Inference

Black-box grammar inference involves inferring the context-free grammar of a programming language by analyzing a set of example strings (e.g., sample programs) without accessing the internal workings of the parser. In black-box grammar inference, the parser is treated as a "black box", where only the inputs (i.e., example strings) and outputs (i.e., parsing results) are visible, while its internal state, rules, or implementation details remain hidden. In contrast, white-box grammar inference provides full access to the parser's internal logic and implementation, allowing for inspection and modification of the parser code to understand how it processes inputs and to directly extract detailed grammar rules. Gray-box grammar inference falls between the two, offering limited internal access that may allow some form of inspection or intermediate state access.

Black-box grammar inference is particularly important in the following scenarios: **Closed-Source Parsers**: Many programming language parsers are proprietary and closed-source, preventing users from accessing or modifying these parsers due to legal or intellectual property restrictions. **Remote Parsers**: Some parsers are available only as remote services, such as online compilers or interpreters, where users can interact with the parser only through API calls or network interfaces, without access to its underlying implementation. **Legacy Systems**: The source code of parsers in legacy systems may be lost or poorly documented, requiring black-box techniques to reconstruct the grammar.

2.3 SOTA - ARVADA and TREEVADA

The state of the art black-box inference of context-free grammar methods include ARVADA [17] and TREEVADA [3]. To construct context-free grammar, ARVADA iteratively perform *bubbling* and *merging* operations on tree representations of the input. We demonstrate how ARVADA infer the context-free grammar of the input "a+3;" and "3 < 3;", as illustrated in Figure 1. ARVADA first construct flat trees from the input, with a single root node t_0 whose children are the characters of the input strings $a + 3 ;$ and $3 < 3 ;$; (Figure 1(a)). The bubbling operation involves taking sequences of sibling nodes in the trees (e.g., a) and adding a layer of indirection (e.g.,

¹<https://github.com/Sinpersrect/kedavra>

t_1) by replacing the sequence with a new node. Then ARVADA decides whether to accept or reject the proposed bubble by checking whether merging the bubbles enables sound generalization of the learned language. When a bubble is merged, ARVADA only accepts it if it expands the language accepted by the generated grammar while preserving the Oracle validity of the strings produced by the generated grammar. Consider our example, a , $+$, 3 , and $<$ are bubbled as t_1 , t_2 , t_3 , and t_4 , respectively (Figure 1(b)). If the string derivable from subtrees can be swapped and still retain a valid grammar, they can be *merged*. ARVADA first attempts to merge bubbles that have similar context. In the above example, ARVADA merges t_2 and t_4 into t_2 because they have similar context (Figure 1(c)). Then, ARVADA iteratively checks any remaining bubbles that can be merged. In this step, t_1 and t_3 are merged into t_1 (Figure 1(d)). These operations are repeated until no remaining bubbled sequence enables a valid merge. Finally, the context-free grammar generated from the example is:

t_0	:	t_1	t_2	t_1	“;
t_1	:	“ a ”		“ 3 ”	
t_2	:	“ $<$ ”		“ $+$ ”	

TREEVADA optimized ARVADA mainly by implementing predefined rules for pairing brackets. In this example, both TREEVADA and ARVADA yield identical results.

3 MOTIVATING EXAMPLE

Let us revisit the generation of the CFG of the example strings: $a+3$; and $3 < 3$; as shown in section 2.3. TREEVADA/ARVADA observes that “ $+$ ” and “ $<$ ” have very similar contexts, as they both precede “ 3 ”. Therefore, TREEVADA/ARVADA attempts to merge “ $+$ ” and “ $<$ ” into t_1 , and then merge “ a ” and “ 3 ” into t_2 .

However, this results in overgeneralization because in TinyC, “ $+$ ” and “ $<$ ” are not interchangeable. For example, multiple instances of “ $<$ ” in a statement like “ $a < a < a$;” are not permitted and considered illegal, whereas multiple “ $+$ ” symbols are allowed, as seen in “ $a+a+a$;”, which is legal.

The method TREEVADA/ARVADA uses to select and merge Bubbles, specifically based on contextual similarity, is not always appropriate in certain scenarios. To address this gap, our proposed method uses an incremental construction strategy for grammar and no longer solely relies on contextual similarity when selecting and merging Bubbles. Instead, it selects Bubbles based on the generalization potential they can bring when merged. In addition to the issue of inaccurate generalization, TREEVADA and ARVADA have several other drawbacks due to inferring the grammar based on all example strings as a whole: (1) low speed, (2) low grammar readability due to complex grammar structures, (3) sensitivity of the inferred grammar to the example strings, and (4) low precision, as detecting overgeneralization in complex grammar is challenging. Our proposed method addresses these drawbacks, leading to improved speed, readability, precision, robustness.

Specifically, the proposed method offers below advantages:

More Accurate Grammar: By evaluating the generalization potential brought by merging Bubbles, our method focuses more on the applicability and extensibility of the generated grammar. This approach enables the generated grammar to better adapt to

different inputs, enhancing the model’s generality and robustness, and ensuring the grammar performs well across various scenarios.

Better Readability: Since our Grammar is constructed incrementally, we can create a much more concise Grammar. This makes it significantly more readable compared to the previous approach.

Faster Processing Speed: With the incremental approach, simpler grammar that have already been processed can be skipped if they appear as part of later, more complex grammar. This approach demonstrated a significant reduction in the inference process time.

4 APPROACH

Figure 2 demonstrates the workflow of the proposed KEDAVRA, which consists of three main steps: **tokenization, data decomposition, and incremental grammar inference**. **Tokenization** takes example strings as input and produce the tokenized sequences that represent syntactically significant units (e.g., keywords, identifiers, operators). Tokenization accelerates subsequent data decomposition and incremental grammar inference. **Data decomposition** takes the tokenized sequence as input and generates a list of decomposed sequences that break down the complex sequences into simpler ones, while collectively preserving all grammatical structures. **Incremental grammar inference** iteratively infers the grammar from the simplest sequences to the most complex ones by decomposition, which remains significantly less complex than previous methods. This approach results in reduced complexity, improved grammar quality, and higher accuracy compared to prior works. We elaborate on the detailed steps in the following subsections.

4.1 Tokenization

Many languages follow fundamental tokenization rules, such as separating identifiers by non-identifier tokens. Before performing grammar inference, KEDAVRA first tokenize the example strings. **This step eliminates the need to repeatedly infer common lexing rules during grammar inference, which is computationally expensive. Consequently, it significantly improves the speed of subsequent grammar inference.**

The proposed tokenization module involves the following steps.

Pre-tokenization: The proposed pre-tokenization applies common lexical rules to segment the example strings into tokens such as identifiers, strings, and numbers. **Special handling is given to whitespace.** Considering that many programming languages are not sensitive to whitespace, treating it as a token can reduce the efficiency of grammar inference (like ARVADA and TREEVADA). Therefore, during pre-tokenization, we also assess whether the target language is sensitive to whitespace. This involves testing whether the inputting of repeated whitespace is accepted by the oracle. If the target programming language does not treat whitespace as significant, we do not consider it as a token in the grammar inference process. An example of program before and after pre-tokenization is demonstrated in Figure 2(a) and Figure 3, respectively.

Token merging: We consider that two tokens can be merged if swapping them in all their occurrences maintains validity (to the oracle). For each pre-tokenized program, we attempt to merge tokens that are swappable. After merging, a token sequence and a token-value dictionary are generated. The dictionary stores the




Figure 2: Workflow of KEDAVRA




Figure 3: Results after pre-tokenization

newly generated tokens and the original tokens that are merged into them. Figure 2(b) shows an example output.

Character-level generalization: The final step in tokenization involves character-level generalization of token values. We compute a character set containing all possible characters corresponding to the token's character type. For instance, if the token's value includes a lowercase letter, we include all lowercase letters in the set. Subsequently, we iteratively replace the original token value with each character from the set and validate the generated token value with the oracle. Furthermore, we experiment with repeating the same type of characters multiple times. For example, if the character is a lowercase letter, we attempt to repeat this letter multiple times to ascertain whether the oracle accepts the input. If the oracle accepts the generated token value, we include it in the character set. Figure 2(c) shows an example of the character-level generalization results.

4.2 Data Decomposition

Previous studies, ARVADA and TREEVADA, have approached grammar inference by analyzing entire example strings. These methods often result in highly complex grammars, especially when the example strings themselves are complex. Such complexity not only slows down the parsing process when these inferred grammars are used in parser construction but also reduces the readability of the generated grammars. The lack of readability adversely affects the user's ability to understand the grammar when it is utilized for creating documentation.

To address this shortcoming, our goal in designing the inference method is to maintain the accuracy of the grammar while making it as simple and readable as possible. To this end, **we propose breaking down complex example strings into simpler components, while collectively preserving all grammatical structures of the original program.** An example of the decomposed sequences and the corresponding input example strings are shown in Figure 2(d) and Figure 2(a), respectively.

Algorithm 1 Data decompose algorithm

Input: tokenized sequence D ; a oracle O
Output: decomposed tokenized sequence

```

procedure DECOMPOSE( $D, O$ )
   $D_{dec} \leftarrow \{\}$ 
  for  $d$  in  $D$  do
    for  $i \leftarrow 1$  to  $|d|$  do
       $d_{dec} \leftarrow$  keep  $d[i]$  and remove other of  $d$ 
       $D_{dec} \leftarrow D_{dec} \cup \{d_{dec}\}$ 
    end for
  end for
  return  $D_{dec}$ 
end procedure

```

We describe the pseudocode of this step in Algorithm 1. In decomposing the data, we follow this method: First, we number all tokens in the sequence - generated from the previous step - from 1 to n . We then preserve the first token and attempt to eliminate as many of the remaining tokens as possible. We assess the reduced token sequence with the oracle to determine its validity; if deemed valid, we continue with the reduction; otherwise, we revert to the original sequence. This process is repeated iteratively until no additional tokens can be removed. We apply this method sequentially for each token from 1 to n . The result of this procedure is n sequences, The i th sequence represents the smallest set of valid token sequences that must contain the i th token. As each token is preserved at least once, this approach ensures that no essential grammar elements are lost while simplifying the sequence.

Figure 2(b) and Figure 2(d) presents an example of the input and output from this step. The comparison clearly shows that the decomposed sequence is simpler than the original input while retaining all grammatical elements from the original input.

4.3 Incremental Grammar Inference

In this section, we detail our methodology for incrementally inferring grammar based on decomposed token sequences. **First, we arrange these sequences by length - from the shortest to the longest - based on the number of tokens.** Our approach incrementally incorporates the least complex sequences into the inference process.

Algorithm 2 outlines our incremental inference approach. **We begin by inferring the grammar from the simplest token sequence generated in the previous step, and continue by iterating over all token sequences derived from the example**

Algorithm 2 Incremental Grammar Inference Algorithm

```

1: Input: tokenized sequences  $D$ ; a oracle  $O$ 
2: Output: grammar
3: procedure INFERGRAMMAR( $D, O$ )
4:    $G \leftarrow \text{CREATEEMPTYGRAMMAR}()$ 
5:   while  $D \neq \emptyset$  do
6:      $T \leftarrow \text{POP SHORTEST}(D)$   $\triangleright$  pop the shortest token sequence in  $D$ 
7:     if  $T \in G$  then
8:       continue
9:     end if
10:     $G \leftarrow \text{ADDPROD}(G, T)$ 
11:    while  $G$  contains production need generalize do
12:       $P \leftarrow \text{CHOOSEPROD}(G)$   $\triangleright$  Choose a production that need generalize.
13:       $B \leftarrow \text{BUBBLING}(G, P, O)$ 
14:       $B_{opt} \leftarrow \arg \max_{i \in B} \text{SCORINGBUBBLESET}(G, i, O)$ 
15:      if  $\text{SCORINGBUBBLESET}(G, B_{opt}, O) = 0$  then  $\triangleright$  If the bubble set's
        score is 0, we skip.
16:        continue
17:      end if
18:       $B_{filtered} \leftarrow \text{ELIMOVERGEN}(G, B_{opt}, O)$ 
19:       $G \leftarrow \text{MERGEBUBBLES}(G, B_{filtered})$ 
20:    end while
21:  end while
22:   $G \leftarrow \text{GENERALIZEREP}(G, O)$ 
23:  return  $G$ 
24: end procedure

```

Algorithm 3 Bubbling algorithm

```

Input: grammar  $G$ ; a production  $P$ ; a oracle  $O$ 
Output:  $\text{bubbleSets}$ 
procedure BUBBLING( $G, P, O$ )
   $\text{bubbleSets} \leftarrow \{\}$ 
  for  $\text{bubble} \in \text{SUBSEQ}(P)$  do  $\triangleright$  Build bubbles by iterate all sub sequence of  $P$ 
     $\text{bubbleSet} \leftarrow \{\text{bubble}\}$ 
    for  $P' \in \text{ALLPRODS}(G)$  do
      for  $\text{bubble}' \in \text{SUBSEQ}(P')$  do
        if  $\text{CHECKSWAP}(\text{bubble}, \text{bubble}', O)$  then  $\triangleright$  Check if two bubbles
          can be swapped.
           $\text{bubbleSet} \leftarrow \text{bubbleSet} \cup \{\text{bubble}'\}$ 
        end if
      end for
    end for
     $\text{bubbleSets} \leftarrow \text{bubbleSets} \cup \{\text{bubbleSet}\}$ 
  end for
return  $\text{bubbleSets}$ 
end procedure

```

strings. The following subsections elaborate on the methodologies implemented in each iteration.

4.3.1 Bubbling. The purpose of bubbling is to identify parts of the grammar that can be generalized.

Algorithm 3 outlines how the bubbling step works. For each production rule P that derived from the input token sequence, we group all its subsequences as bubbles B_p . Then we also group all the subsequences of the grammar's production rules as bubbles B_{all} . Here, we employ the TREEVADA's approach to handle parentheses. We discard bubbles with unmatched brackets, such as those containing a left bracket without a corresponding right bracket.

Then we iterate through B_p and assess whether each bubble can be swapped with bubble in B_{all} . We define two bubbles as *swappable* if the grammar, after the swap, can still be validated by the target oracle. For each bubble B_i in B_p , we create a bubble set $\{B_i, B_j, B_{j+1}, \dots, B_m\}$, where B_j, B_{j+1}, \dots, B_m is all bubbles in B_{all} that can be swapped with B_i .

The generated bubble sets are then fed into the next step.

4.3.2 Choose the Most Generalizable Bubble Set. Algorithm 4 outlines how to calculate the bubble set's score for choosing the most generalizable bubble set.

We choose the bubble set that has the highest generalization score, indicating the potential to produce the most generalizations. Consider bubble set $\{B_i, B_j, B_{j+1}, \dots, B_m\}$. We first swap B_i and B_j , and get two sample string from respectively position. then swap B_i and B_{j+1} until B_i and B_m are swapped. The number of samples that result in a new generalization can serve as the generalization score. This score reflects the extent of generalization achievable by the bubble set.

A new generalization is considered to be created if we swap two bubbles so that the resulting grammar contains a string accepted by the oracle but not included in the grammar before the merge.

We then choose the bubble set with the highest score for further processing in the next step.

4.3.3 Eliminating Over-generalization. The bubble set $\{B_i, B_j, B_{j+1}, \dots, B_m\}$ selected in the previous step only examines whether B_i are swappable with other bubbles (*i.e.*, B_j, B_{j+1}, \dots, B_m). However, we did not check if the bubbles generated from the existing grammar (*i.e.*, B_{j+1}, \dots, B_m) are swappable. Merging these unswappable bubbles may introduce over-generalization, defined as the grammar containing strings that cannot be accepted by the oracle. In this step, we aim to select a subset of bubbles that maximize generalizability without causing over-generalization. Algorithm 5 outlines our method for this step.

First, we try to find a minimal subset with full generalization: we remove all bubbles in the subsets that do not affect the generalization score. The generalization score is calculated as follows: First we collect all the generalized samples from the grammar that merging the bubbles in section 4.3.2. Then we evaluate how many of these samples are in the grammar after merging the bubble set (the bubble set we need to estimate generalization score). This rate is the generalization score. This process yields minimal complete generalization subsets. **Based on this subset, we seek to find a subset that 1) does not over-generalize, and 2) has the highest generalization score.**

To achieve this, we enumerate all subsets of the current bubble set, then select all subsets that do not cause over-generalization.

Algorithm 4 Calculate bubble set score

```

1: Input: grammar  $G$ ; a bubble set  $S$ ; a oracle  $O$ 
2: Output: bubble set score
3: procedure SCORINGBUBBLESET( $G, S, O$ )
4:    $B_{first} \leftarrow \text{CHOOSEFIRSTBUBBLE}(S)$ 
5:    $S_{remain} \leftarrow S \setminus \{B_{first}\}$ 
6:    $\text{samples} \leftarrow \{\}$ 
7:   for  $B \in S_{remain}$  do  $\triangleright$  Swap first bubble with each of others, and record the
     samples
8:      $\text{samples} \leftarrow \text{samples} \cup \text{SWAPBUBBLE}(B_{first}, B)$   $\triangleright$  We swap the two
       bubble, then we get two samples
9:   end for
10:   $\text{samples}_{gen} \leftarrow \{i \in \text{samples} \mid \text{CHECKGEN}(i, G, O)\}$   $\triangleright$  Check how many
    of the swapped samples can generalize.
11:  return  $|\text{samples}_{gen}|$ 
12: end procedure
13: procedure CHECKGEN( $\text{Sample}, G, O$ )  $\triangleright$  We consider a sample that not in grammar
    but can accepted by oracle as generalized.
14:  return  $\text{Sample} \notin G \wedge \text{CHECK}(\text{Sample}, O)$ 
15: end procedure

```

Algorithm 5 Eliminating Over-generalization

```

1: Input: grammar  $G$ ; a bubble set  $S$ ; a oracle  $O$ 
2: Output: bubble set without over-generalization
3: procedure ELIMOVERGEN( $G, S, O$ )
4:    $G' \leftarrow \text{MERGEBUBBLE}(G, S)$ 
5:    $\text{Samples} \leftarrow \{i \in \text{SAMPLE}(G') \mid \text{CHECKGEN}(i, G, O)\}$  ► We create
     samples of  $G'$  that indicate the full generalization, it can be used to estimate the
     generalization score of  $S$ 's subset
6:    $S_{\min} \leftarrow S$ 
7:   for all  $b \in S$  do
8:     if GENSCORE( $S_{\min} \setminus \{b\}, G, \text{Samples}$ ) = 1 then ► If we remove a bubble
       does not reduce the generalization score, we remove it
9:        $S_{\min} \leftarrow S_{\min} \setminus \{b\}$ 
10:    end if
11:  end for
12:   $\text{subsets} \leftarrow \{S_{\text{sub}} \subseteq S_{\min} \mid \text{CHECKOVERGEN}(S_{\text{sub}})\}$  ► We choose all
     subset of  $S_{\min}$  that does not cause over-generalization
13:  return arg max $_{i \in \text{subsets}}$  GENSCORE( $i, G, \text{Samples}$ )
14: end procedure
15: procedure GENSCORE( $B, G, \text{Samples}$ )
16:    $G' \leftarrow \text{MERGEBUBBLES}(G, B)$ 
17:   return  $|\{i \in \text{Samples} \mid i \in G'\}| / |\text{Samples}|$  ► Check how many
     samples are in the grammar after merging Bubbles.
18: end procedure

```

From these, we choose the subset with the highest generalization score.

We would like to note that KEDAVRA exhibits non-determinism, primarily during the process for eliminating over-generalization. In this phase, we utilize a sampling method to select grammar samples for estimating the generalization score, which introduces non-determinism into our approach.

4.3.4 Merging and Grammar Simplification. Merging involves uniting all bubbles into a single non-terminal. We observed that the merged grammar may introduce redundancy. To mitigate this redundancy and accelerate the inference process, we further simplify the generated grammar. For example, if a non-terminal only appears in one production rule and this production rule does not contain any other non-terminal or terminal, we merge this non-terminal with the production rules it produces and so on. It is important to note that these simplify rules were observed during the experiments and may not encompass all possible redundancies. However, new rules can be added if additional cases are identified. Notably, simplifying the grammar speeds up the inference process without impacting accuracy.

4.3.5 Generalize Rep and Expansion of Terminals. After iterating through all sequences generated from the example strings (as described in Section 4.2), we identify the repeatable bubbles within the grammar. This is achieved by sample generated with repeated bubbles against the oracle. If the sample generated by the repeated bubbles is accepted by the oracle, the bubble is marked as repeatable.

At this point, we obtain a context-free grammar inferred from the example strings. The final step involves expanding each terminal in the grammar to a larger character class, based on the character-level generalization table produced in Section 4.2. This ensures that the grammar can accept tokens of the same type that were not present in the example strings but are still valid. For instance, expanding $t_1 \rightarrow a \mid b \mid c$ to include all lowercase letters.

Table 1: Example strings S with their avg char size; # = nr. programs;

	R0		R1		R2		R5	
	#	avg	#	avg	#	avg	#	avg
arith	17	2.3	n/a					
fol	36	14.5	n/a					
math	62	5.5	n/a					
json	71	3.9	30	11.7	30	8.6	n/a	
lisp	26	2.5	30	79.2	30	24.8	n/a	
turtle	33	7.7	35	41.1	35	25.6	n/a	
while	10	15.5	30	171.4	30	217	n/a	
xml	40	11.6	20	27.8	20	27.5	n/a	
curl	25	20.7	25	22.1	25	22.0	n/a	
tinyc	25	80.5	25	96.2	25	86.4	10	514
minic	10	107	n/a					

5 EVALUATION

We designed a novel algorithm with the aim of inferring as many high-quality grammars as possible. To evaluate whether this goal has been fulfilled, we applied KEDAVRA to 8 micro benchmarks and 3 macro benchmarks, comparing it to the most cutting-edge tools. We then provide an in-depth analysis of the comparison results to answer the following five research questions.

RQ1 Grammar quality: Does KEDAVRA infer better grammars than TREEVADA and ARVADA on the same dataset?

RQ2 Runtime: Is the runtime of KEDAVRA lower than TREEVADA and ARVADA on the same dataset?

RQ3 Readability: How compact are the inferred grammars?

RQ4 The impact of different sampling methods on grammar accuracy.

5.1 Experimental Setup

To investigate the effectiveness and efficiency of KEDAVRA in grammar inference, we applied our tool on a slightly modified benchmark dataset², which includes a total of 11 grammars, each of which contains 1k test programs. To this end, we include all the original datasets from ARVADA³ and TREEVADA⁴. We removed Nodejs from the original dataset because the Nodejs dataset's oracle was incorrectly implemented, leading to erroneous outputs. For instance, the oracle would consider " $()=}$ " as syntactically correct. In reality, this input is incorrect, but the oracle used Nodejs's `-check` feature and checked the output for 'SyntaxError'. However, when the input contains multiple errors, other errors can obscure the 'SyntaxError'. This results in the oracle incorrectly considering an erroneous input as syntactically correct. For example, " $()=}$ " would trigger a 'ReferenceError' instead of a 'SyntaxError', causing the oracle to wrongly deem it as syntactically correct.

In addition, to compensate for the reduced dataset resulting from the removal of the Node.js dataset and to further evaluate KEDAVRA's capability with sophisticated grammar features, we provided a larger and more complex C program called "minic" which includes function calls, variable declarations, and function declarations. Additionally, we utilized KEDAVRA's sampling algorithm to generate test and training data for this program. At the same time,

²<https://github.com/Sinpersrect/kedavra>

³<https://github.com/neil-kulkarni/arvada>

⁴<https://github.com/rifatarefin/treevada>

Table 2: Rerun ARVADA (left), rerun TREEVADA (middle) and KEDAVRA(right) on the same example strings; ARVADA/KEDAVRA values are average over 10 runs; f1 = F1 score; \pm = standard deviation; bold = KEDAVRA better than or equal to ARVADA& TREEVADA

	ARVADA			TREEVADA			KEDAVRA		
	p	r	f1	p	r	f1	p	r	f1
arith	1.0 \pm .0	1.0 \pm .0	1.0 \pm .0	1	1	1	1.0\pm.0	1.0\pm.0	1.0\pm.0
math	.97 \pm .03	.75 \pm .2	.83 \pm .13	1	.47	.64	.94 \pm .03	.92\pm.1	.93\pm.07
fol	.99 \pm .03	.73 \pm .31	.8 \pm .23	1	.53	.69	.99 \pm .04	.55 \pm .18	.69 \pm .14
json	.91 \pm .09	.97 \pm .09	.93 \pm .06	.98	.94	.96	1.0\pm.0	.97 \pm .02	.99\pm.01
lisp	.97 \pm .06	.38 \pm .26	.5 \pm .19	.7	.73	.71	1.0\pm.0	1.0\pm.0	1.0\pm.0
turtle	1.0 \pm .0	1.0 \pm .0	1.0 \pm .0	1	1	1	1.0\pm.0	1.0\pm.0	1.0\pm.0
while	1.0 \pm .0	.75 \pm .23	.84 \pm .15	1	.24	.38	1.0\pm.0	1.0\pm.0	1.0\pm.0
xml	1.0 \pm .0	.96 \pm .1	.98 \pm .06	1	1	1	1.0\pm.0	1.0\pm.0	1.0\pm.0
curl	.58 \pm .08	.92 \pm .02	.71 \pm .06	.8	.71	.75	1.0\pm.0	.11 \pm .12	.18 \pm .19
tinyc	.55 \pm .24	.76 \pm .36	.6 \pm .29	.74	.9	.81	1.0\pm.0	1.0\pm.0	1.0\pm.0
minic	.65 \pm .32	.29 \pm .17	.32 \pm .11	.63	.62	.63	1.0\pm.0	.48 \pm .0	.65 \pm .0

we also leveraged KEDAVRA’s sample method to generate test and training data for this program. Hence, the overall dataset for the experiment contains 25 train datasets correspond to 11 grammar. Table 1 shows the size and average char size of each train dataset. Our experiment runs on a Linux server with Intel(R) Xeon(R) Silver 4210R CPU @ 2.40GHz and 128GB RAM.

To facilitate easy comparison, we conducted the experiments using TREEVADA’s Docker image and reused its blackbox parser. It’s worth noting that when assessing the precision of the generated grammar, we not only attempted to replicate the previous sampling strategy as closely as possible but also constructed a new sampling strategy: We limit the usage of each production rule to no more than 10 times. This approach addresses the issue of the previous sampling method’s inability to detect deeper grammar errors.

5.2 RQ1: Grammar Quality

Our first research question concerns the quality of grammar generated by KEDAVRA and how it compares to existing tools. In this work, like many other grammar inference approaches, we prioritize precision, recall and F1 score to evaluate the grammar quality. Specifically, these three metrics are calculated as follows:

- **Precision:** We sampled 1,000 cases from the inferred grammar and then assessed how many of these were accepted by the oracle to calculate the precision,
- **Recall:** Based on a set of 1,000 "golden" test programs, we calculate the recall by determining how many of the inferred grammar rules appear in these programs,
- **F1 score:** F1 score is the harmonic mean of precision and recall.

Result. Our experimental results (on R0) are presented in table 2, which shows the performance (*i.e.*, precision, recall, and F1 score) of KEDAVRA compared to two state-of-the-art (SOTA) tools targeting the same problem of grammar inference. Here, we chose ARVADA [17] and TREEVADA [3] as baseline because they are recognized as the most cutting-edge tools [3] and have been made publicly available in the community.

Overall, as highlighted in table 2, among the 11 grammar datasets, the F1 score of KEDAVRA (*i.e.*, *arith*, *math*, *json*, *lisp*, *turtle*, *while*, *xml*, *tinyc* and *minic*, respectively on average) is higher than or equals to that of ARVADA and TREEVADA, further indicating that our tool performs better on most of the datasets. For example, KEDAVRA achieves high scores on *arith*, *lisp*, *turtle*, *while*, *xml*, and *tinyc* with a precision, recall, and F1 score of 1.0 \pm 0.0. Even for the largest program input dataset, Minic, KEDAVRA still maintains good




Figure 4: avg F1 score of 10 runs of ARVADA, TREEVADA and KEDAVRA on each dataset (R0, R1, R2, R5). Note that the horizontal bars in each of the sub-figures are manually added as a reference to better visualize the fluctuations in the F1 scores of the inference algorithm across different datasets.

performance, with a high precision of 1.0 \pm 0.0 and an F1 score of 0.65 \pm 0.0. This high score experimentally shows the effectiveness of KEDAVRA in generating grammar with high quality. The outliers are *fol* and *curl*, where KEDAVRA’s f1 score (*i.e.*, 0.69 and 0.18) is below that of ARVADA (*i.e.*, 0.8 and 0.71) and TREEVADA (*i.e.*, 0.69 and 0.75). This is because in *fol*, we did not merge the optimal bubbles when inferring the grammar, which caused us to lose generalization in the next step, resulting in a lower F1 score compared to ARVADA and TREEVADA. In *curl*, our tokenization algorithm does not fit *curl*’s grammar well, as it is not similar to a programming language. Consequently, the poor tokenization result leads to low recall in *curl*. Further discussion of the performance on *curl* can be found in Section 6. On the other hand, KEDAVRA’s recall in *minic* is lower than in TREEVADA, which can be mainly explained by the limitations of grammar loss, where the grammar of function calls in *minic* might be lost in data decomposition because we treat a function call as a combination of an identifier and an expression. During the process of data decomposition, only one of them will be retained. Overall, despite of these imprecise results, our approach is effective in generating better grammar compared to SOTA tools.

In addition to evaluating the grammar quality, it remains unknown whether KEDAVRA can consistently maintain consistent performance across diverse datasets. To this end, we further analyse the F1 scores on different datasets (*i.e.*, R0, R1, R2, R5) with the same grammar, as presented in Figure 4. This figure displays the F1 score of three tools across seven grammar datasets, using the distance of the bars from the horizontal axis to represent the deviation of the F1 scores. Note that the horizontal bars in each of the sub-figures are manually added as a reference to better visualize the fluctuations in the F1 scores of the inference algorithm across different datasets. Intuitively, we find that, our approach shows relatively small variations (*i.e.*, KEDAVRA has the shortest distance from the horizontal axis) in the learned results, whereas previous methods are heavily influenced by dataset characteristics. Across various grammars such as Lisp, Turtle, TinyC, XML, and While, our approach consistently achieves F1 score of 1 on different datasets. In contrast, TREEVADA and ARVADA exhibits significant f1 score variations across different datasets of *json*, *lisp*, *turtle*, *while*, *xml* and *tinyc*. This indicates that KEDAVRA achieves consistent performance across different datasets corresponding to the same grammar. After carefully comparing the design of KEDAVRA with the other two approaches, we attribute the greater stability of KEDAVRA’s F1 score on most datasets to the following factor: We infer grammar from

decomposed data rather than the original data. Starting with an empty grammar, we incrementally construct the grammar based on the decomposed data. Since the smaller segments in the decomposed data are highly similar and contain minimal grammar, the resulting grammar remains relatively stable.

Answer to RQ1: KEDAVRA outperforms state-of-the-art grammar inference tools, TREEVADA and ARVADA, by achieving higher precision, recall, and F1 scores on most datasets. Experimental evidence also underscores the consistency of KEDAVRA in grammar inference.

5.3 RQ2: Resources

While effectiveness measures the quality of the inferred grammar, efficiency also serves as a vital factor as it assesses the operation in a timely and resource-effective manner. Thus, in this research question, we further examine the efficiency of KEDAVRA by reporting the time cost and computational resources. Specifically, for the sake of easy comparison, we report the following metrics in alignment with the SOTA work: **Runtime:** The time cost in executing the tool. **Oracle Time:** ARVADA, TREEVADA and KEDAVRA rely on external parser for oracle validation. Thus we measure the time cost during the oracle checking process. **Oracle Calls:** We also calculate the number of oracle calls. **Peak Memory:** We use the Linux 'time' command to measure peak memory usage during grammar inference.

Table 3 shows the performance of resource utilization for ARVADA, TREEVADA and KEDAVRA on R1&R5 (c-500 is R5 dataset of *tinyc*). Overall, KEDAVRA takes 1.2 kiloseconds on average to process an program to do grammar inference, which is faster than that of the TREEVADA (*i.e.*, on average, 3.4 kiloseconds to process a program) and the ARVADA (*i.e.*, on average, 7.4 kiloseconds to process a program). As experimentally demonstrated by Arefin [3], increasing the grammar quality will likely reduce efficiency as it requires trade-offs when handling the frequency of the validations. It explains why KEDAVRA takes more time on *fol*, *math*, *json* and *xml* dataset as KEDAVRA takes extensive validation steps during the grammar inference process to ensure correctness. These validation steps require the use of an oracle to verify the precision of the inferred grammar rules, resulting in a slight increase of time cost. However, we argue that it is inevitable to make such trade-offs between effectiveness and efficiency. In addition, the fact that the time difference between KEDAVRA and TREEVADA is relatively small (*i.e.*, *fol*, *math*, *json* and *xml*), their runtime difference is only a few hundred seconds, which is completely acceptable. On more complex datasets, our method can still maintain a fast runtime due to the data decomposition step. it suggests that KEDAVRA is applicable for use, despite its slightly longer execution time for some grammars, given its superior effectiveness and lower time overhead.

Moreover, it's worth mentioning that KEDAVRA spends considerably less time than ARVADA and TREEVADA on *lisp*, *while*, *tinyc*, *minic* and *c-500* datasets. Upon scrutinizing the time expenditure of each sub-process, we ascertain that this is primarily due to KEDAVRA's innovative data decomposition process, which deliberately circumvents redundant learning of identical grammars present in

the original data. In essence, once KEDAVRA establishes a grammar rule, there is no need for further re-learning in subsequent processes, significantly enhancing efficiency.

For memory consumption, we observed a significant reduction in memory usage for *tinyc* and *c-500*, while for other grammars, the difference is small (at most a twofold difference). The reduced memory usage for *tinyc* and *c-500* is primarily due to the data decomposition, which greatly reduces the runtime. Since KEDAVRA's memory usage mainly stems from caching Oracle results, the reduction in runtime naturally leads to lower memory consumption.

Answer to RQ2: KEDAVRA has an average runtime of 1.2 kiloseconds and memory usage of 0.1 GB for each dataset. Compared to ARVADA and TREEVADA, there is no significant increase in memory usage, while the average runtime significantly decreases.

5.4 RQ3:Readability

ARVADA, TREEVADA do not try to simplify their inferred grammars; they merely export the grammar when they cannot identify any further generalization steps. But KEDAVRA will simplify the inferred grammar after merge bubbles to speed up subsequent processing. Given a use-case involving human consumption like program understanding, it is worthwhile to investigate the readability of the generated grammar. Another relevant use-case is that overly complex grammars often contain numerous ambiguities, which can lead to difficulties for semantic exploration due to the presence of ambiguity. To this end, in this research question, we leverage the following metrics to evaluate readability:

- **Grammar Size:** We calculate the count of unique non-terminals and terminals, along with the number of productions and the length of each production (*i.e.*, the length of the right-hand-side sequence of terminals and non-terminals). The size of the grammar is then determined by summing the lengths of all its productions.
- **Parse Time & Memory:** We assess the total time and peak memory usage needed to parse the 1k "golden" test programs using a parser generated from a grammar inferred by ARVADA, TREEVADA or KEDAVRA.

The results are presented in Table 4, indicating among the most datasets (highlighted in bold), our approach achieves significantly higher readability score compared to previous methods. Specifically, not only achieving higher precision, recall and F1 score, but also KEDAVRA's scores on NT, A and S are smaller than that of the other two tools. For example, regarding rule lengths, KEDAVRA's grammars are smaller for 9 out of 12 languages. The most substantial difference is observed in the *tinyc*, *minic*, *c-500*, *lisp*, *while* languages, where KEDAVRA's grammar is less than half the size of the ARVADA/TREEVADA grammar. The exceptions are *fol*, *math*, and *xml*, where for the *math* and *xml* languages, KEDAVRA's F1 score is higher despite a slight difference in grammar size (*e.g.*, *xml*'s 178 vs. 156). In terms of *fol*, although KEDAVRA's grammar is larger, we argue that the overhead is relatively small (189 vs. 159), which is acceptable.

Table 3: Average ARVADA & TREEVADA results over 10 runs on R1& R5(c-500 is R5 of tiny); t = runtime; t_O = oracle time; q = queries sent to oracle; m = peak memory usage; \pm = standard deviation; bold = KEDAVRA better than TREEVADA & ARVADA

	ARVADA				TREEVADA				KEDAVRA			
	t [ks]	t_O [ks]	q [k]	m [GB]	t [ks]	t_O [ks]	q [k]	m [GB]	t [ks]	t_O [ks]	q [k]	m [GB]
arith	0±0	0±0	5±0	.02±0	0±0	.00±0	0.40	0.02	0±0	0±0	6±0	.08±0
fol	.33±0	.2±0	28.3±3.3	.03±0	.15±0	.12±0	16.30	0.03	.67±.2	.52±.1	58.0±5.5	.08±0
math	.06±0	.04±0	6.9±.6	.02±0	.05±0	.04±0	6.30	0.02	.64±.5	.17±.1	20.9±7.4	.11±.1
json	.17±0	.11±0	13.9±2.6	.05±0	.05±0	.04±0	6.30	0.02	.08±.1	.05±0	7.1±.5	.13±0
lisp	10.4±2.2	2.1±1.7	40.7±25.6	18.3±22.3	1.6±.1	.96±0	52.60	0.06	.27±0	.25±0	16±.1	.12±.1
turtle	.83±.3	.25±.1	29.1±8.4	.06±0	.17±0	.11±0	16	0.04	.15±0	.13±0	19±.1	.03±0
while	11.7±3.5	2.8±1.0	49.6±14.6	.88±.1	3.3±.1	.39±0	31.60	0.13	1.62±.1	1.5±.1	123.8±6.5	.28±.2
xml	.38±.1	.15±0	19.3±2.1	.05±0	.22±0	.09±0	10.80	0.08	.94±.1	.82±.1	88±8.6	.10±0
curl	.44±0	.4±0	22.3±1.3	.05±0	.38±0	.35±0	21.10	0.04	.31±0	.29±0	17.5±.2	.04±0
tinyc	12.2±3.9	5.1±2.2	86.5±25.7	1.2±1.0	5.5±.2	1.9±.1	114.30	0.23	.16±0	.09±0	21.8±1.1	.04±0
minic	15.3±5.2	13±5.1	42.6±17.9	1.4±.8	17.1±.0	15.9±0	66.70	0.12	8.9±1.3	8.7±1.2	32.7±5.2	.13±.1
c-500	37.0±13.0	15.7±7.5	137.4±48.8	3.0±3.6	11.7±1.0	3.1±.2	145.90	0.42	.25±.1	.13±0	23.4±2.3	.05±0

Table 4: Grammar size and parse performance on R1 & R5(c-500 is R5 of tiny): Averages for grammars inferred in 10 runs; NT/T = unique (non-) terminals; A = rules (alternatives); l(A) = avg. rule length; S = sum of rule lengths; tP = time to parse 1k “golden” test programs; mP = peak memory while parsing; bold = KEDAVRA better than ARVADA&TREEVADA.

	ARVADA							TREEVADA							KEDAVRA						
	NT	A	l(A)	S	T	t_P [ks]	m_P	NT	A	l(A)	S	T	t_P [ks]	m_P	NT	A	l(A)	S	T	t_P [ks]	m_P
arith	6	31	1.2	38	16	.05	.07	6	30	1.2	36	16	.05	.07	6	29	1.2	34	16	.11	.08
fol	18	98	1.5	145	80	.04	.04	20	108	1.5	159	80	.04	.03	11	145	1.3	189	80	.02	.03
math	22	116	1.3	150	71	.07	.06	19	119	1.4	166	71	.02	.03	18	181	1.1	205	72	.27	.10
json	31	158	1.4	217	74	.02	.02	17	178	1.2	208	82	.01	.02	13	114	1.2	137	74	.01	.02
lisp	22	104	9.6	738	40	.27	.05	18	95	1.7	157	40	.14	.04	5	69	1.2	80	40	.05	.03
turtle	28	122	1.6	192	67	.04	.03	15	95	1.3	124	67	.03	.03	11	86	1.2	104	67	.01	.02
while	36	91	2.1	188	18	.26	.20	24	55	2	113	18	.11	.09	7	17	3.3	55	18	.03	.04
xml	19	92	1.5	141	58	.02	.02	9	78	2	156	58	.02	.02	7	120	1.5	178	58	.01	.02
curl	23	181	1.4	253	95	.87	.07	18	135	1.4	186	82	.09	.03	20	121	1.2	140	80	.05	.03
tinyc	56	221	2.2	460	52	.50	.07	30	150	1.8	264	51	.39	.04	10	57	1.6	89	49	.03	.02
minic	37	214	2.4	507	82	.20	.06	39	209	1.8	376	82	.15	.04	19	74	1.7	124	55	.01	.02
c-500	59	240	2.4	521	52	.77	.11	41	188	1.8	345	51	.24	.05	9	57	1.6	90	49	.05	.02

In addition, in terms of time, the fact that KEDAVRA spends much less time on most language datasets demonstrates that the larger grammars of ARVADA and TREEVADA do not improve parsing performance. On the contrary, for 9 out of 12 experiments, KEDAVRA’s parse time is less than half of TREEVADA’s and ARVADA’s. However, there are several outliers (*i.e.*, arith, math, and json), where KEDAVRA has a higher F1 score despite a slight difference in time consumption. Consider memory consumption, In half of the datasets (*i.e.*, c-500, minic, tinyc, while, turtle and lisp), our memory usage is lower than TREEVADA and ARVADA, while in the other half it is higher. Overall, their memory usage does not differ significantly.

We conducted an ablation study to further evaluate the impact of each component of KEDAVRA on the readability of the inferred grammar. Since some components form the foundation of KEDAVRA and cannot function independently (*e.g.*, KEDAVRA would not operate if either the Data Decomposition or Incremental Inference modules were removed), we opted to integrate each component into TREEVADA rather than removing them from KEDAVRA. This approach allowed us to assess how these additions could enhance the readability of TREEVADA’s outputs. Consequently, the experiments were designed around the following three scenarios: TREEVADA with the Data Decomposition module, TREEVADA with the Grammar Simplification module, and TREEVADA with both the Data Decomposition and Grammar Simplification modules. It is important to note that Incremental Inference is the core framework for grammar inference and not a modular component; therefore, it cannot be added to TREEVADA for experimentation. The original TREEVADA was used as the baseline for comparison.

Table 5: Grammar Readability on R0: Averages for grammars inferred in 10 runs; TREEVADA + D = TREEVADA with Data Decomposition; TREEVADA + S = TREEVADA with Grammar Simplification; TREEVADA + D&S = TREEVADA with Data Decomposition and Grammar Simplification. NT = unique non-terminals; A = rules (alternatives); S = sum of rule lengths.

	TREEVADA			TREEVADA + D			TREEVADA + S			TREEVADA + D&S		
	NT	A	S	NT	A	S	NT	A	S	NT	A	S
arith	6	30	36	6	28	34	5	29	35	5	27	33
fol	20	108	159	18	100	148	18	101	146	17	98	145
math	19	119	166	15	82	115	14	95	121	13	72	95
json	15	178	210	11	108	137	10	104	132	10	103	129
lisp	10	54	73	4	33	44	8	50	66	4	33	44
turtle	11	84	106	8	75	89	11	80	97	7	74	88
while	8	18	49	8	18	49	8	18	49	8	18	49
xml	10	82	125	10	81	125	9	73	108	9	73	103
curl	19	141	196	7	82	99	19	136	186	7	81	96
tinyc	35	164	293	18	105	177	31	130	218	18	87	133
minic	39	209	376	23	183	264	36	185	326	21	181	262

As shown in Table 5, both Data Decomposition and Grammar Simplification contribute to improved readability, as indicated by the reductions in NT, A, and S. Specifically, Grammar Simplification reduces NT by 12%, A by 16%, and S by 17%, while Data Decomposition achieves reductions of 33%, 25%, and 28% for these metrics, respectively. This indicates that Data Decomposition has a more significant impact on enhancing readability. When combined, Data Decomposition and Grammar Simplification further reduce NT, A, and S by 38%, 29%, and 34%, respectively, demonstrating their effectiveness in increasing the readability of the inferred grammar.

Answer to RQ3: KEDAVRA not only achieves higher precision, recall, and F1 score, but also outperforms TREEVADA and ARVADA in readability for the majority of languages.




Figure 5: Precision of ARVADA, TREEVADA and KEDAVRA runs on different sample methods (A = ARVADA’s sample algorithm, T = TREEVADA’s sample algorithm, LPP10 = LimitPerProd10). Note that the horizontal bars in each of the sub-figures are manually added as a reference to better visualize the fluctuations in the precision values of the inference algorithm across different sampling algorithms.

5.5 RQ4 : Sampling Method

Apart from the effectiveness, efficiency, and readability of KEDAVRA, it is also worth investigating the sampling algorithms as they can have a huge effect on precision calculation. Actually, in our experiments, we discovered that the previous sampling methods (*i.e.*, those provided by ARVADA and TREEVADA) could not detect deeper grammatical errors because they limit the maximum recursion depth to avoid generating excessively long samples. Specifically, The ARVADA sampling algorithm rejects samples longer than 300. When sampling grammars, if the recursion depth exceeds 5, it tends to use production rules that only contain terminal symbols. The TREEVADA sampling algorithm is similar to the ARVADA sampling algorithm but does not limit sample length.

To avoid such shortcomings, We designed a more reliable algorithm that limits excessively long samples in a different way, named LimitPerProd10 (LPP10), as follows:

LimitPerProd10: We limit the usage of each production rule to no more than 10 times. If a production rule is used more than 10 times, we replace the current non-terminal symbol with a pre-computed string corresponding to that non-terminal symbol (to limit the maximum length of the sample).

We then applies the three sampling algorithm on a dataset for each grammar (R0 for arith,fol,math,minic. R1 for others). The results are presented in Figure 5, demonstrating the precision of ARVADA, TREEVADA, and KEDAVRA using three different sampling algorithms. Note that the horizontal bars in each of the sub-figures are manually added as a reference to better visualize the fluctuations in the precision values of the inference algorithm across different sampling algorithms. In all the grammar datasets, KEDAVRA has the least variation (*i.e.*, KEDAVRA has the shortest distance from the horizontal bar), whereas the precision of all three tools is heavily influenced by the sampling algorithm. Here, taking tinyc as an example, ARVADA’s minimum and maximum precision values of

different sampling algorithm are respectively 0.11 and 0.79. TREEVADA’s minimum and maximum precision values of different sampling algorithm are respectively 0.22 and 0.88. However, KEDAVRA’s precision consistently high with minimal variation, demonstrating excellent stability.

To further explain why LPP10 outperforms the ARVADA and TREEVADA sampling algorithms, we manually checked some cases and discovered that ARVADA and TREEVADA have high precision rates with their own sampling algorithms but perform much worse with the LPP10 algorithm. This is because both the ARVADA and TREEVADA sampling algorithms overlook deep grammars, which are, however, exposed by the LPP10 algorithm. Therefore, the results from the LPP10 sampling algorithm better reflect reality.

Here, we proposed an example to show the differences:

```

n0 : n1
n1 : n2
n2 : n3
n3 : n4
n4 : n5
n5 : n6 | "1"
n6 : "2"

```

In the above example, assuming the oracle accepts only “1” and not “2”. When sampling grammar, reaching n_5 results in maximum recursion depth for ARVADA and TREEVADA algorithms. In this case, these algorithms tend to choose production rules with terminal symbols, picking “1” instead of n_6 , yielding a precision of 1. However, under the LPP10 sampling algorithm, each production rule’s usage count is limited; for example, once $n_5 \rightarrow n_6$ is used more than 10 times, it won’t be selected again. This allows us to still have the probability of selecting n_6 when reaching n_5 , enabling us to detect deeper errors.

In summary, we proposed a more reliable sampling algorithm that can be used for evaluation by other grammar inference tools.

Answer to RQ4: We have identified an issue with the sampling algorithms of TREEVADA and ARVADA, which fail to detect deep errors in complex grammars. To address this, we propose a new sampling algorithm that can more accurately evaluate grammar accuracy.

6 DISCUSSION AND FUTURE WORKS

The availability of example strings. Current black-box context-free grammar inference approaches, including GLADE [5], ARVADA [17], TREEVADA [3], and our KEDAVRA, all require both an oracle (*i.e.*, a black-box parser) and the availability of example strings. These example strings must conform to the grammar of the target language, and the quality of the inferred grammar is significantly influenced by the quality of these strings—for example, whether the strings adequately cover most of the language’s grammar. However, in real-world settings, example strings are not always readily available, which restricts the practical application of existing black-box grammar inference methods. A notable attempt to circumvent the need for example strings was made by REINAM [34], which employed

an industrial symbolic execution engine to generate initial example strings before conducting grammar inference. This approach, however, is feasible only in white-box scenarios, where the internal details of the oracle are accessible, and is unsuitable for black-box grammar inference. Looking ahead, bridging this gap constitutes a critical area for future research. Developing methodologies capable of initiating grammar inference without the need for pre-existing example strings would greatly enhance the adaptability and application of black-box grammar inference across various domains. Such advancements could significantly broaden the scope and utility of grammar inference techniques in environments constrained by the lack of data accessibility.

The impact of the quality of the example strings. The quality of the example strings - for example, whether the strings adequately cover most of the language's grammar - inherently impacts the quality of the inferred grammar. KEDAVRA addresses this challenge through a data decomposition step, which significantly reduces the impact of example quality on the grammar generated, compared with ARVADA and TREEVADA. In KEDAVRA, the decomposed strings derived from different example strings tend to be very similar. This effect is illustrated in Figure 4, where the F1 score variation of the grammar inferred by KEDAVRA across different datasets is much smaller than that observed with ARVADA and TREEVADA.

Performance on *curl* and conventional programming languages. KEDAVRA utilizes a general tokenizer that incorporates lexical rules effective across many programming languages, recognizing substrings matching the regular expression “[a-zA-Z][0-9a-zA-Z]*” as identifiers and “[0-9]+” as numeric tokens in most programming languages like Python, Go, Java, C/C++, etc.. This design choice optimizes KEDAVRA for conventional programming languages where these patterns are prevalent. However, these assumptions are less effective for languages like *curl*, or those with identifiers that include characters beyond alphanumeric and underscores, such as *Perl*, *Ruby*, *JavaScript*, and *Scala*. KEDAVRA's poor results with *curl* arise from incorrect handling of URLs. For example, the tokenization step inappropriately splits the URL “http://1.g.i10/?” into “http, :, /, /, 1, ,, g, ,, i10, /, /, ?” with tokens as “t0: // t1 t2 t2 t2 t2 // t3”, whereas each character after “http://” should be treated uniformly (i.e., “t0: // t1 t1 t1 t1 t1 t1 t1 t1 t1”). This demonstrates that while KEDAVRA is effective with conventional programming languages whose grammars follow typical identifier rules, it struggles with non-standard token systems due to its foundational assumptions. Addressing these specific limitations and exploring potential optimizations will be part of our future work.

Scalability. The scalability of KEDAVRA is mainly determined by the tokenization processing. Tokenization time increases quadratically with the size of the sample set due to two key steps: substring splitting, which takes $O(n)$, and merging mergeable tokens, which takes $O(n^2)$. In contrast, incremental inference, which uses decomposed strings for grammar inference, is less impacted because different samples may produce the same decomposed strings. We would like to highlight that KEDAVRA is less influenced by the sample set size compared to ARVADA and TREEVADA. For instance, in our experiments with the *tiny*c datasets (Table 3), specifically R1 and R5 — where R5 has a larger sample set — the runtime difference for ARVADA/TREEVADA is 2-3 times greater between the datasets. In contrast, the runtime increase for KEDAVRA from R1 to R5 is only 1.5

times, demonstrating its relative scalability under increased input sizes compared to the state of the art.

7 RELATED WORKS

Black-box Grammar Inference. GLADE [5] initiated black-box grammar inference, but its performance declines with highly recursive grammars. ARVADA [17] advanced this by better handling recursive grammars, while TREEVADA [3] improved the process by optimizing for bracket structures and recursive grammar application, allowing the inference of deterministic grammars. REINAM [34], differing from CFG, infers a target program's PCFG[14], first deriving an initial CFG with GLADE, then enhancing it through reinforcement learning to adjust probabilities within the PCFG.

Deep Learning. Research [6, 29, 36] has assessed the limitations of RNN, in learning CFG. Despite deep learning's inability to directly derive CFG, Yellin [35] developed a method to extract CFG from RNN. However, results reveal substantial shortcomings of LSTM compared to ARVADA, likely due to the absence of *active learning* [2] which is utilized in ARVADA, TREEVADA, and KEDAVRA.

Grey-box Grammar Inference. Grey-box inference leverages partial information from an oracle without full access to its source code. GRIMOIRE [7] demonstrates this by deriving grammar-like information from oracle's coverage data during fuzz testing, albeit not generating a direct CFG.

White-box Grammar Inference. This approach, exemplified by Lin [19] and AUTOGRAM [13], relies on full access to oracle's source code, using dynamic analysis and dynamic taint analysis respectively to trace inputs and deduce CFGs. Its application limited to scenarios where source code access is permissible.

Semantics Exploration. Semantic exploration has been researched by many people [16, 32, 38], such as tools like ISLa [32], focuses on inferring semantics within grammar by defining input constraints, either manually or mined from samples. This exploration is dependent on existing grammar, suggesting that an integrated approach with grammar derivation could significantly enhance the generation of high-quality inputs for black-box programs.

8 CONCLUSIONS

Inferring context-free grammars from black-box environments poses substantial challenges due to the restricted availability of example programs. Current leading methods, ARVADA and TREEVADA, utilize heuristic strategies to derive grammar rules from basic parse structures while examining a range of generalization sequences. These techniques, however, tend to produce grammars of lower quality and readability, a consequence of processing complete example strings which increases complexity and computational time. In response, we developed a novel methodology, KEDAVRA, that decomposes example strings into smaller segments for incremental grammar inference. Our method consistently outperforms existing approaches in terms of grammar precision, recall, computational efficiency, and readability, as verified by our empirical studies.

9 ACKNOWLEDGMENT

This work was supported by the Key Laboratory of Computing Power Network and Information Security, Ministry of Education under Grant No.2023ZD034.

REFERENCES

- [1] Ziyad Alsaeed and Michal Young. 2023. Finding Short Slow Inputs Faster with Grammar-Based Search. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 1068–1079.
- [2] Dana Angluin. 1981. A note on the number of queries needed to identify regular languages. *Information and Control* 51, 1 (1981), 76–87.
- [3] Mohammad Rifat Arefin, Suraj Shetiya, Zili Wang, and Christoph Csallner. 2024. Fast Deterministic Black-box Context-free Grammar Inference. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–12.
- [4] Cornelius Aschermann, Tommaso Frassetto, Thorsten Holz, Patrick Jauernig, Ahmad-Reza Sadeghi, and Daniel Teuchert. 2019. NAUTILUS: Fishing for Deep Bugs with Grammars. In *NDSS*.
- [5] Osbert Bastani, Rahul Sharma, Alex Aiken, and Percy Liang. 2017. Synthesizing program input grammars. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, 95–110.
- [6] Jean-Philippe Bernardy. 2018. Can recurrent neural networks learn nested recursion? *Linguistic Issues in Language Technology* 16 (2018).
- [7] Timo Blazytko, Matthew Bishop, Cornelius Aschermann, et al. 2019. GRI-MOIRE: Synthesizing structure while fuzzing. In *28th USENIX Security Symposium (USENIX Security 19)*. 1985–2002.
- [8] Juan Caballero, Heng Yin, Zhenkai Liang, and Dawn Song. 2007. Polyglot: Automatic extraction of protocol message format using dynamic binary analysis. In *Proceedings of the 14th ACM conference on Computer and communications security*. 317–329.
- [9] A. Cremers and S. Ginsburg. 1975. Context-free grammar forms. *J. Comput. System Sci.* 11, 1 (1975), 86–117.
- [10] Patrice Godefroid, Adam Kiezun, and Michael Y Levin. 2008. Grammar-based whitebox fuzzing. In *Proceedings of the 29th ACM SIGPLAN conference on programming language design and implementation*. 206–215.
- [11] Rahul Gopinath and Andreas Zeller. 2019. Building fast fuzzers. *arXiv preprint arXiv:1911.07707* (2019).
- [12] James Gosling, Bill Joy, Guy Steele, Gilad Bracha, Alex Buckley, Daniel Smith, and Gavin Bierman. 2022. *The Java Language Specification: Java SE 19 Edition*. Oracle.
- [13] Matthias Hörschle and Andreas Zeller. 2016. Mining input grammars from dynamic taints. In *Proc. 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*. ACM, 720–725.
- [14] Frederick Jelinek, John D. Lafferty, and Robert L. Mercer. 1992. *Basic methods of probabilistic context free grammars*. Springer Berlin Heidelberg.
- [15] Tae-Woong Kim, Tae-Gong Kim, and Jai-Hyun Seu. 2013. Specification and automated detection of code smells using OCL. *International Journal of Software Engineering and Its Applications* 7, 4 (2013), 35–44.
- [16] Željko Kovačević, Marjan Mernik, Matej Ravber, et al. 2020. From grammar inference to semantic inference—An evolutionary approach. *Mathematics* 8, 5 (2020), 816. <https://doi.org/10.3390/math8050816>
- [17] Neil Kulkarni, Caroline Lemieux, and Koushik Sen. 2021. Learning highly recursive input grammars. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 456–467.
- [18] Thomas Köppe et al. 2022. *Working Draft, Standard for Programming Language C++*. Technical Report N4917. ISO/IEC.
- [19] Zhiqiang Lin and Xiangyu Zhang. 2008. Deriving input syntactic structure from execution. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 83–93.
- [20] Benoît Marchal. 2002. *XML by Example*. Que Publishing.
- [21] Leon Moonen. 2001. Generating robust parsers using island grammars. In *Proceedings eighth working conference on reverse engineering*. IEEE, 13–22.
- [22] Csaba Nagy and Anthony Cleve. 2017. A static code smell detector for SQL queries embedded in Java code. In *2017 IEEE 17th International Working Conference on Source Code Analysis and Manipulation (SCAM)*. IEEE, 147–152.
- [23] John Narayan, Sandeep Shukla, and Charles Clancy. 2015. A survey of automatic protocol reverse engineering tools. *Comput. Surveys* 48, 3 (2015), 1–36. <https://doi.org/10.1145/2732254>
- [24] Hoang Lam Nguyen, Nebras Nassar, Timo Kehrler, and Lars Grunke. 2020. Mo-fuzz: A fuzzer suite for testing model-driven software engineering tools. In *Proc. 35th IEEE/ACM International Conference on Automated Software Engineering*. 1103–1115.
- [25] Yusuke Oda, Hiroyuki Fudaba, Graham Neubig, Hideaki Hata, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2015. Learning to generate pseudo-code from source code using statistical machine translation. In *Proc. 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 574–584.
- [26] Yasubumi Sakakibara. 1990. Learning context-free grammars from structural data in polynomial time. *Theoretical Computer Science* 76, 2-3 (1990), 223–242.
- [27] Yasubumi Sakakibara. 1992. Efficient learning of context-free grammars from positive structural examples. *Information and Computation* 97, 1 (1992), 23–60.
- [28] Yasubumi Sakakibara and Hiroshi Muramatsu. 2000. Learning context-free grammars from partially structured examples. In *Grammatical Inference: Algorithms and Applications: 5th International Colloquium, ICGI 2000, Lisbon, Portugal, September 11-13, 2000. Proceedings*. Springer Berlin Heidelberg, 229–240.
- [29] Luzi Sennhauser and Robert C Berwick. 2018. Evaluating the ability of LSTMs to learn context-free grammars. *arXiv preprint arXiv:1811.02611* (2018).
- [30] Ben Smith. 2015. *Beginning JSON*. Apress.
- [31] Prashast Srivastava and Mathias Payer. 2021. Gramatron: Effective grammar-aware fuzzing. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*. 244–256.
- [32] D. Steinhöfel and A. Zeller. 2022. Input invariants. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 583–594.
- [33] Junjie Wang, Bihuan Chen, Lei Wei, and Yang Liu. 2019. Superion: Grammar-aware greybox fuzzing. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 724–735.
- [34] Z. Wu, E. Johnson, W. Yang, et al. 2019. REINAM: reinforcement learning for input-grammar inference. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 488–498.
- [35] Daniel M Yellin and Gail Weiss. 2021. Synthesizing context-free grammars from recurrent neural networks. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 351–369.
- [36] Xiang Yu, Ngoc Thang Vu, and Jonas Kuhn. 2019. Learning the Dyck language with attention-based Seq2Seq models. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*. 138–146.
- [37] Shu yu Guo, Michael Ficarra, and Kevin Gibbons. 2022. *ECMAScript 2022 Language Specification (13th ed.)*. Technical Report ECMA-262. Ecma International.
- [38] C. Zhou, Q. Zhang, L. Guo, et al. 2023. Towards better semantics exploration for browser fuzzing. In *Proceedings of the ACM on Programming Languages*, Vol. 7. 604–631.