

SYNShine: Improved Fixing of Syntax Errors

Toufique Ahmed^{ID}, Noah Rose Ledesma, and Premkumar Devanbu

Abstract—Novice programmers struggle with the complex syntax of modern programming languages like Java, and make lot of syntax errors. The diagnostic syntax error messages from compilers and IDEs are sometimes useful, but often the messages are cryptic and puzzling. Novices could be helped, and instructors' time saved, by automated repair suggestions when dealing with syntax errors. Large samples of novice errors and fixes are now available, offering the possibility of data-driven machine-learning approaches to help novices fix syntax errors. Current machine-learning approaches do a reasonable job fixing syntax errors in shorter programs, but don't work as well even for moderately longer programs. We introduce SYNShine, a machine-learning based tool that substantially improves on the state-of-the-art, by learning to use compiler diagnostics, employing a very large neural model that leverages unsupervised pre-training, and relying on multi-label classification rather than autoregressive synthesis to generate the (repaired) output. We describe SYNShine's architecture in detail, and provide a detailed evaluation. We have built SYNShine into a free, open-source version of Visual Studio Code (VSCode); we make all our source code and models freely available.

Index Terms—Deep learning, program repair, naturalness

1 INTRODUCTION

SYNTAX errors are easy to make, and will cause compiles to fail. The challenges posed by syntax errors to novices have been known for a long time [1]. More recent studies have documented the challenges faced by novices in various languages [2], [3], [4]. Novices make a wide range of syntax mistakes [4], some of which are quite subtle; time that might otherwise be spent on useful pedagogy on problem-solving and logic is spent helping novices deal with such errors. Unfortunately the error messages provided by compilers are often not helpful; novices struggle to interpret the messages, and sometimes even experts do! [5]. Consider for example, the real program example in Fig. 1, where a novice student just replaced a “`”` with an “`x`” on line 8. None of the big 4 IDEs (VSCode, IntelliJ, BlueJ, or Eclipse) provide a direct diagnostic for this very understandable error. A lot of time can be spent on such errors [6], and researchers have called out for more attention to help novices [5] deal with errors, specifically syntax errors. While *semantic errors* (bug-patching) have received quite some attention, syntax errors have attracted less interest.

The possibility of collecting novice error data, and the emergence of high-capacity, highly configurable deep learning models, has raised the possibility of designing models that can automatically fix errors, and training them using novice data. This approach is very attractive for several reasons: a) Automated repair of syntax errors is helpful to novices, and

saves instructors' time. b) Traditional approaches to automatically finding & fixing syntax errors require hand-coding fairly complex parser logic. c) Automatically learning models to fix errors is an approach that promises to be language-agnostic, as long as sufficient data is available. d) Learning fixing strategies from samples representative of novice errors promises to yield models that perform well on the most common mistakes that novices make. Several recent approaches to this problem have emerged, which are all arguably language-agnostic.

DeepFix [7] used sequence to sequence encoder/decoder models (with roots in language translation) to fix all sorts of errors in C, while Santos et al. [8] used language models to repair just syntax errors in Java (and thus is closer to our work). All of the existing approaches take an erroneous program as input, and attempt to fix them. DeepFix (which uses an RNN-based Seq2Seq approach) works less well on longer programs, since RNNs struggle with long-range dependencies. Santos et al. faced similar challenges. More recently, Ahmed et al. [9] trained a “lenient” parser using synthetic data. Ahmed et al. use a 2 stage approach: a first (“BLOCKFix”) repair the nesting structure, and the second (“FRAGFix”) to repair individual statements; we refer to their tool in this paper as BF + FF, to indicate their two stages. Their approach improves over both DeepFix and Santos et al. especially for longer programs. Syntax errors in longer programs (longer than 200-300 tokens) are challenging for automated repair, because *locating* the error is difficult. All the above approaches ignore an important source of information that could be of great value: *the error from the compiler!* Compiler warnings include a lot of useful information: including often the line number where the error occurred, the tokens involved in the error, and the nature of the error. This information could be used by neural model to better locate and repair the error. Yasunaga et al. [10] utilize C compiler warnings with a graph-based self-supervised approach and outperform DeepFix in fixing compiler errors. However, compiler wanrnings have not been applied to any approach that is specifically designed for Java programs. Our model also uses compiler warnings,

• The authors are with the Department of Computer Science, University of California, Davis, CA 95616 USA. E-mail: {tfahmed, roseledesma, ptdevanbu}@ucdavis.edu.

Manuscript received 12 October 2021; revised 20 September 2022; accepted 29 September 2022. Date of publication 10 October 2022; date of current version 18 April 2023.

This work was supported by the U.S. National Science Foundation under Grants 1414172 and 2107592.

(Corresponding author: Toufique Ahmed.)

Recommended for acceptance by A. Mesbah.

Digital Object Identifier no. 10.1109/TSE.2022.3212635

```

1 import java.util.Scanner;
2 public class Multiplication
3 {
4     public static void main(String[] args){
5         Scanner sc = new Scanner(System.in);
6         int a = sc.nextInt();
7         int b = sc.nextInt();
8         int res = a * b;
9         System.out.println("The result is: " + res);
10    }
11 }

```

Fig. 1. Incorrect novice code sample.

but our performance remains robust as the programs' length increases.

In addition, existing approaches have not adequately exploited the tremendous capacity of current DL models to learn (without direct supervision) the statistics of *very large* amounts of unlabeled sequential data. Modern pre-training approaches such as RoBERTa can ingest vast corpora of sequential data (e.g., a billion tokens from GitHub-hosted code) and learn the patterns of syntax, identifier usage patterns, arithmetic expressions, method call patterns *etc.* These patterns are automatically learned and represented as high-dimensional vector embeddings of tokens, without requiring any human effort to label the data. These embeddings, however, have been shown to substantially improve performance when used as pre-set embeddings in other networks that can be “fine-tuned” with smaller amounts of human-labeled data.

In this paper, by using the diagnostics from a compiler, and exploiting the ability to pre-train embeddings with high capacity RoBERTa model, we build a tool, SYNShine, which improves substantially on the state-of-the-art in automated syntax repair in JAVA. We make the following contributions:

- 1) We utilize compiler diagnostics from javac, as well as unsupervised pre-training to achieve substantial improvements, to implement a 3-stage syntax error repair tool, which can fix as much as 75% of programs with single errors in the Blackbox dataset. This substantially improves upon prior work in the area of “JAVA” syntax error repair.
- 2) When generating fixes, we rely on multi-label classification, rather than autoregressive synthesis, to simplify the task of generating the repair.
- 3) We evaluate the contributions of the different stages of our tool, and also the value of pre-training, and the use of javac.
- 4) We evaluate the diversity of repairs that SYNShine can perform; we also dig into the cases where it appears to fail.
- 5) We have built SYNShine based repair into the widely used, freely available, open-source Visual Studio Code (VSCode) tool, and made all our software and data available to the extent allowable under legal requirements¹.

Note: Most of the novice code correction approaches are designed for C including DeepFix [7], [10], [12], [13]. Some recent works [10], [13] outperform DeepFix in fixing C

compiler errors. They all take the complete program as input and evaluate it on the DeepFix dataset with smaller sequences (up to 450 tokens). Ahmed et al. have already shown that models taking complete program sequences tend to fail more often for longer programs [9]. Unlike Blackbox, DeepFix dataset does not have erroneous and fixed program pairs. That prevents us from comparing the model's performance with the human-produced fixed versions. We train DeepFix model on our JAVA dataset because DeepFix uses the simplest inductive bias: sequence of program tokens and does not depend on any language-specific compiler. Several other approaches [10], [13] are both *compiler-* and *language-* dependent, so they are not comparable with our approach. Furthermore, we are able to accept complete programs, of longer length than earlier approaches, and provide fixes leveraging both pre-training as well as compiler errors.

2 BACKGROUND & MOTIVATION

Problem-solving, motivation & engagement, and difficulties in learning the syntax of programming language are three fundamental challenges in introductory programming courses [14]. The dropout and failure rates are still high in introductory programming courses even after applying advanced methods and tools [15], [16]. Helping novices with programming syntax can prevent novices to get demotivated [14] at the beginning of the learning process. In this paper, we aim to help novice programmers by automatically suggesting repairs for syntax errors. Consider the program in Fig. 1, which is an actual example our dataset of novice programs with errors [11]. Note the use of “x” instead of “*” on line 8. Many school maths texts use “x” for multiply, so this is an understandable error.

In an introductory programming course, a novice may make this error by force of habit, and then find it quite challenging to fix the problem. Most popular IDEs (Eclipse, IntelliJ, Visual Studio Code) have trouble fixing this; however, our approach, which feeds a javac-based error diagnostic, into a multi-stage repair engine that combines unsupervised pre-training, with fine-tuning, can resolve this.

Researchers have been interested in compiler diagnostics or syntax error messages for over half a century [17]. Barik et al. reported [18] that the difficulties programmers face while reading or understanding error messages are comparable to the difficulty of reading source code. Understanding Java error messages is quite challenging for two reasons; i) the same error produces different diagnostics depending on the context, and ii) the compiler may produce the same diagnostic for different errors [18]. Though prior works [19], [20] addressed fixing errors in novice programs, DeepFix [7] was the first to apply deep learning to fix errors. DeepFix considers code repair as Neural Machine Translation (NMT) and uses an encoder-decoder based deep learning model to fix errors in C programs. Though initially aimed at semantic bugs, the approach also works for syntax errors. This approach was limited by the use of RNN (recurrent neural network) seq2seq models—the RNN architecture is challenged by longer inputs, and outputs; also since the back-propagation through time (for the recursive elements) is not easily parallelized, it's challenging to exploit

1. Blackbox data is distributed under U.K. Laws. Please contact creators [11] for details.

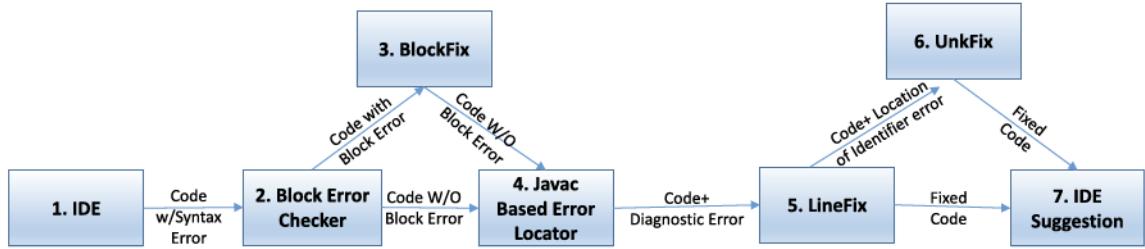


Fig. 2. Overall architecture of the SYNSHINE tool.

larger datasets and additional processors. These became nagging problems in NLP; initial efforts with basic attention mechanisms [21] were supplanted by powerful multilayer models with multiple attention heads to avoid recursive elements altogether [22], yielding high-capacity, eminently parallelizable *transformer* models. Certain errors, such as the ones relating to block nesting, statement delimitation (with “;”) etc. involve long-range syntax dependencies, and require attending to very long contexts, which transformers can do better; still, even these models fail when the dependencies become much longer.

Ahmed et al. [9], developed BF + FF, using a multi-layer, multi-head transformer approach, to address the limitations of traditional seq2seq models. In addition, BF + FF used a two-stage pipeline, with the first stage addressing long-range block nesting errors, even ones beyond the range of transformers (BLOCKFIX) and the second stage addressing shorter-range errors (FRAGFIX). Using the Blackbox [11] dataset, they demonstrated that their approach substantially improved over prior work on the same dataset [8] (which used language models). BF + FF had important limitations, noted in their paper; it didn’t take advantage of error localization and diagnosis provided by compilers; it also didn’t effectively address errors in identifiers. Indeed, none of the existing approaches dealt effectively with identifiers, since they had to limit vocabulary. Deep learning models are challenged by large vocabularies, which require very large embedding and softmax layers. (See [23] details). We use BPE [23] to address this issue.

By addressing these limitations, we were able to achieve very substantial improvements on the state of the art for fixing JAVA programs. Ahmed et al.² and Gupta et al.³ provided extensive source-available replication packages which enabled us to provide a detailed comparison (See §4).

3 METHODOLOGY

Previous work had various limitations: longer programs were difficult to repair; error messages from compilers were not used; vocabulary limitations in DeepFix and design choices in BF + FF limited the ability to address errors in identifier usage. SYNSHINE directly addresses these issues, and achieves substantial improvements. We use a multi-stage pipeline which incorporates the Java programming language compiler (javac), along with three learned DL neural networks (DNN). The first DNN model is directly based on the BLOCKFIX stage provided by BF + FF; this resolves (the potentially long-range dependent) nesting

errors in the program. In the second stage, SYNSHINE departs from BF + FF. BF+FF uses the fixed nesting structure from BLOCKFIX to split the program into lines, and then just tries to fix *every line*; this leads to a lot of incorrect fixes. Deepfix and Santos et al. also try to fix the entire program. The second stage (LINEFIX) in SYNSHINE uses the line-location of the error, as detected by the standard javac compiler, together with the actual error message, and generates relevant fixes for delimiters, operators, and keywords; it also flags potential locations for errors in identifier usage; these locations are sent to the third & final stage, UNKFix. The UNKFix DNN model uses a Roberta-MLM to correct any identifiers that flagged as potentially wrong by LINEFIX.

3.1 Overall Architecture

Fig. 2 shows the architecture of our approach. When the IDE flags an error (step 1) we first pass the program through a block-nesting error checker (2), which is a simple pushdown automaton, that checks the program’s nesting structure. If block-related issue is found, it’s sent from (2) to BLOCKFIX (3) a transformer model (as provided in the open-source BF + FF implementation [9]) for repair. In either case, the code, hopefully now free of block-nesting errors, is sent to step 4, where we try to locate the erroneous line using javac. We identify the line that javac associates with the syntax error, and pass it on to LINEFIX (step 5) with the error message. In some cases, LINEFIX can fix it directly; in others, it passes a token position to UNKFix (6), primarily to fix errors in identifier usage. Finally, the fixed code is returned as a suggestion to the IDE (7).

We separate the line-level repairs into LINEFIX and UNKFix to eke out more functions out of deep-learning model capacity. LINEFIX outputs one of 154 possible editing commands, to insert/delete/substitute delimiters, keywords, operators, or identifiers. We limit its output vocabulary to 154. This limitation improves performance, but results in more “unknown” fixes, as described further below (§3.4). These unknowns are resolved by the final DNN model, UNKFix. UNKFix uses a high-capacity masked-language model to suggest a fix (usually an identifier being renamed or inserted) given a location. In combination, these elements allow us to substantially surpass the state-of-the-art.

3.2 Javac Errors: Promises and Perils

While novices often find compiler error messages unhelpful [5], our own experience suggests that they do help experienced developers! This suggests that with sufficient training data, machine-learning models could learn something about how to fix syntax errors, from compiler syntax-error diagnostics. Older machine-learning-based approaches had not

2. <https://zenodo.org/record/4420845>

3. <https://bitbucket.org/iiscseal/deepfix/src/master/>

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on September 23, 2023 at 14:10:12 UTC from IEEE Xplore. Restrictions apply.

<pre> public class PrintVsPrintln { public static void (String args []) { System.out.print(3); System.out.println(4+5); } } </pre> <p>(a) Program with syntax error</p>	<pre> a_50.java:3: error: <identifier> expected public static void (String args []) ^ a_50.java:5: error: illegal start of expression System.out.print(3+); ^ 2 errors </pre> <p>(b) javac error message before fixing</p>
	<pre> a_50.java:1: error: class PrintVsPrintln is public, should be declared in a file named PrintVsPrintln.java public class PrintVsPrintln ^ 1 error </pre> <p>(c) javac error message after fixing</p>

Fig. 3. Locating erroneous line using `javac`.

leveraged these diagnostics [7], [8], [9]. Recently, DrRepair [10] uses these diagnostics for fixing C programs; SYNShine also uses them.

`javac` flags syntactically incorrect programs with diagnostic errors; though the messages *are not precise*, they are sometimes useful. Fig. 3a presents an actual novice program with two syntactic errors (missing “*main*” and unwanted operator “ $+$ ”). The `javac` compiler reports those two errors for the given program 3(b). Although these error messages are unhelpful, `javac` does in this case finger the actual lines with errors. Line-level syntax error localization can be helpful, if the program is long. DeepFix, for example, can not fix longer programs; it relies on seq2seq translation methods, and so has trouble with inputs longer than a few 100’s of tokens. BF + FF resolves this problem by trying to fix every line in the program using its FRAGFix second stage; this approach does induce a fair number of false positives. `javac` promises more accurate location, which could reduce this risk.

There is a potential issue with using `javac`, arising mainly from the constraints of our novice error Blackbox dataset. `javac` generates some error categories which cannot be fixed by editing the program directly. These errors arise for example, from file-naming conventions and incomplete typing environments. For example, class name & file-name mismatch errors, and missing class definition errors are shown in Fig. 3c. The Blackbox dataset (also used by Santos et al. [8] and Ahmed et al. [9]) only includes programs with errors and their associated fix; it does not include the complete programming environment. SYNShine only deals with errors that can be fixed by *directly editing the JAVASOURCE*; we ignore the others. This is a decision also made by all the other papers that deal with syntax error correction [7], [8], [9]; we do, however, make use of compiler diagnostics for JAVA, and do manage to fix a much larger portion of the errors in the Blackbox dataset than prior work, as seen in Table 2. Therefore, to remove the errors we don’t consider from our training set, we simply wrote a wrapper around `javac`, to retain just those errors that can be fixed by editing the source⁴. However, it is important to note here that these “unfixable” errors in our dataset are counted in the denominator when we report our final success rate; in other words, these errors excluded from

training are counted against SYNShine and other tools as failures, and are not ignored in our reported performance.

3.3 Recovering Block Structure: BLOCKFix

Errors involving imbalanced curly braces are prevalent in novice programs, and are hard to resolve because of the long distance between the pair of braces. Ahmed et al. [9] report that block nesting errors consist of around 20-25% of all syntactic errors in novice programs [9]. They incorporate a component, BLOCKFix, for fixing block-nesting errors. BLOCKFix uses a transformer-based machine-translation model to locate & fix block-nesting errors; the translation model is trained on synthetic data with artificially generated nesting errors, and the corresponding fix. It works with an abstracted version of the code without statements, identifiers, and types to fix errors in nesting structure. In SYNShine, we simply adopt the BLOCKFix component from the implementation made available by Ahmed et al.’s replication package.

Ahmed et al. abstracted out all the identifiers, constants, expressions, and delimiters, retraining just the curly braces and keywords (see Fig. 4). They then introduce structure-related syntax corruptions, by adding or dropping the curly braces at randomly chosen positions; and then teaching the model to recover the original abstracted version from the corrupted model. BLOCKFix model learns to fix such errors by training on many such abstracted, corrupted pairs. After fixing the nesting error, the abstracted tokens are replaced with the original ones, and the program is passed to the following stages for further processing.

We found that `javac` works quite well in localizing the error (at least the buggy line and finding the line is sufficient for our approach) *if the program is free of nesting errors*. This is why we apply BLOCKFix, before running `javac` to localize and diagnose the error.

3.4 Fixing Line Error: LINEFix

LINEFix uses a RoBERTa based pre-training + fine-tuning approach. RoBERTa derives from BERT, which uses unlabeled text data to pre-train deep bidirectional representations of text by jointly conditioning on both left and right context in all layers of a deep transformer model [24] to perform simple, self-supervised tasks like filling in masked tokens. This model and training method effectively captures

4. Most common ignored errors relate to “file and class name mismatch” and “undeclared identifiers”.

the statistics of token co-occurrences in very large corpora within the layers of the transformer model. This pre-trained model learns excellent vector representations of code patterns in the higher layers of the transformer; these learned vector representations can be “fine-tuned” with just one additional output layer for specific tasks, and achieves state-of-the-art performance. For pre-training, BERT uses two tasks: fill in masked out tokens using the context (also known as Masked language modeling, or “MLM”) and predict the next sentence given the previous one (the “NSP” task). Liu et al.’s RoBERTa (Robustly Optimized BERT Pre-training Approach) dominates BERT’s performance [25]. Liu et al. drop the NSP objective but dynamically change the masking pattern used in the MLM of BERT models.

Pre-training + fine-tuning also works very well indeed for code. One can gather millions of unlabeled code tokens from open-source projects, conduct pre-training, and then fine-tune the model with a limited amount of labeled data to achieve state-of-the-art performance in different software engineering applications [26], [27], [28], [29] (albeit not yet for code syntax repair). Since we are working on novice code correction and our objective does not involve any relation between two programs, such as Question Answering (QA) and Natural Language Inference (NLI), training on NSP is not beneficial. Furthermore, using a dynamic masking pattern to the training data helps the model achieve better performance in downstream tasks. Therefore, We use RoBERTa for pre-training and fine-tuning of the model.

Why Pre-Training? As explained in the papers on BERT [24] and RoBERTa [25], for natural language, and the very recent, but rapidly growing body of literature using pre-training for code [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], pre-training is a way to exploit enormous volumes of data in a self-supervised fashion to learn the statistics of token sequences, and capture patterns in a position-dependent vector notation. For our purposes, these pre-trained models are automatically ingesting patterns of syntax and identifier usage from vast quantities of source code (around a billion tokens) and bringing all this knowledge implicitly to bear to the task of fixing errors in syntax and identifier usage.

Pre-Training To generate the dataset for pre-training, we collected 5000 most starred Java projects from GitHub (since our end-goal is to correct Java syntax errors). We tokenized the files, yielding 1.2 billion tokens for the pre-training. For the MLM pre-training over code, we randomly select 15% of tokens, and replace with a unique token mask. The loss here is the cross-entropy of the original masked token. Of the 15% selected tokens, 80% are replaced with a specific marker `mask`, 10% are left unchanged, and a randomly selected token replaces the remaining 10%. This training method follows the standard RoBERTa protocol.

The architecture is as shown in Fig. 5. The main RoBERTa model is in the central grey box, labeled “RoBERTa” in Figs. 5a and 5b. The left side is the architecture when RoBERTa is being pre-trained; the last layer on top is the MLM, implemented as a softmax layer taking the RoBERTa embeddings as input, and produces an output token. The entire model is trained using cross-entropy loss. Our RoBERTa architecture consists of 12 attention layers, 768 hidden dimensions, and 12 self-attention heads in each layer. We

```
public class Main {
    public static void main(String[] args) {
        int x = 7;
        int y = 8;
        int sum = x + y;
        System.out.println(sum);
    }
}
```

(a) Original function

```
public class simple_name { public static void
simple_name paren_expression { expression
expression expression }}
```

(b) Abstracted version

Fig. 4. Abstracting source code for recovering block structure.

applied Byte Level BPE (Byte Pair Encoding) tokenizer [23] limiting the sub-token vocabulary size to 25K.

We trained the MLM model using cross-entropy loss on two NVIDIA Titan RTX GPUs for five epochs with a batch size of 44 sequences and learning rate $5e - 5$. When pre-training completed, our MLM model achieved a final loss corresponding to a perplexity of 1.46, (cross-entropy 0.546 bits) which is rather low; RoBERTa for natural language yields final losses around 3.68-4.0 perplexity (1.88 to 2 bits).

Fine-Tuning The fine-tuning step here is to train LINEFix, a model that accepts an incorrect input line from a novice program, (the line flagged by javac as containing a syntax error) together with the text of the error itself, and then generates a set of locations and edit commands, using multi-label classification layers, as explained below.

For fine-tuning and then for evaluation, we used realistic novice programs with syntax errors and human-produced fixed versions. We used the exact dataset used by Santos et al. [8] and Ahmed et al. [9] from the Blackbox [11] repository. This dataset contains 1.7M pairs, of erroneous and fixed programs. Both Santos et al. and Ahmed et al. primarily report their performance on programs with a single token error because a single edit can fix a large fraction of the programs (around 57%). Therefore, for a fair comparison, we also initially focused our evaluation on single token errors and broke down our performance by token-length, as done by Ahmed et al. We selected a test set of 100K samples, with samples stratified by length, from the full dataset for the evaluation. We divided the test dataset into ten token-length ranges (lengths of 1-100, 101-200, ..., and 900-1000 tokens), with each range having around 10K examples. We prepare our fine-tuning dataset from the remaining examples.

Since BLOCKFix handles long-range block-nesting errors, the LINEFix stage is focused on those errors unrelated to nesting. We discarded the programs with imbalanced curly braces from the training set, and after tokenization, we found around 540K examples to train the model. We used javac (discussed in Section 3.2) to localize the error. The input to the model then is the buggy line indicated by javac, appended with a special separator token (denoted `<SEP>`) followed by the error message from javac. Altogether, the maximum input is 150 sub-tokens, which captures virtually all the input lines flagged as erroneous in our dataset. From this, the pre-trained RoBERTa model calculates positional embeddings for each subtoken; however, as with many RoBERTa-based classification tasks, we use just the embedding of the first token.

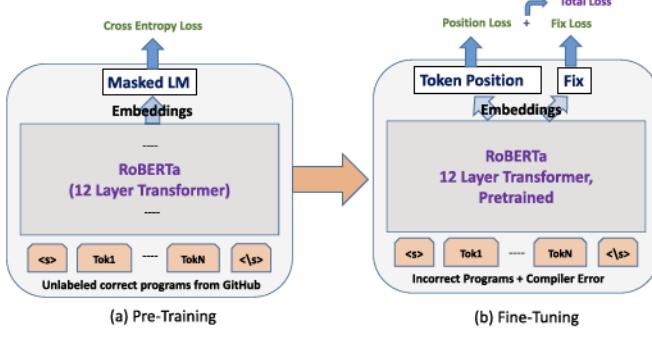


Fig. 5. Pre-training and fine-tuning using RoBERTa.

The desired output is the matching edits required to create the fixed version, as explained next.

To make a complete fix, the model should produce one or more locations, and one or more “fix”, *viz* edit commands. The fix has two parts: i) the *type* of fix (insertion, deletion, or substitute?) ii) the *content* of the fix (is it a specific keyword, delimiter, or any other token?). When the *type* is a deletion, there is no *content* required: if the model identifies the buggy token at position x and recommends deletion, we just drop that token. For substitute operation, if the location is x and the edit command is *substitute* \rightarrow y , we will replace the token at position x with the token y . For insertion, if the command for position x is *insert* \rightarrow y , we will add the suggested y token at the $x + 1$ position. For insertion at the start of the line, we use a special token. For example, consider the following buggy line from Fig. 3a.

```
public static void (String args[])
```

To fix this missing “main”, LINEFix should output the location “3” and the fix “*insert* → *unk*” (“main” is an identifier). This “unk” will be converted to “main” with another model. We will discuss it in Section 3.5.

Our model’s final layer consists of two distinct multi-label classification output layers, one which outputs one or more locations, another which outputs one or more fixes. The input to both these output layers, as explained above, is the RoBERTa embedding of the first token of the input. From this input, the two separate multi-label classification output layers calculate the position(s), and fix(es). Since most (99%) of the erroneous lines are 100 tokens or less, we output one or more positions (1-100) from the first output layer, and, from the second output layer we generate one or more of 154 distinct possible fixes. We remind the reader that a multi-label classification task involves generating an output vector of class probabilities, where the classes are non-exclusive. A single input might generate one or more class labels. In our case, we take all class labels in the output vector scoring above 0.5 as an assigned label. If none of the classes are assigned a probability above 0.5, we just take the highest probability class label. In almost all cases, we have only one fix per line, so one position and one edit command are expected; however, in rare cases, more than one position *and* more than one edit command could be generated. In the former case, we just apply the edit command at that position; in the latter case, which occurs very rarely, we try all combinations and return the first edit combination that compiles. A somewhat more common case (for example

with multiple missing delimiters, like ")", we get one edit command like *insert* →) and multiple locations, in which case, we just apply the same edit at all locations.

There are reasons for our choice of multi-label classification, rather than simply synthesizing the fixed output. Prior approaches [7], [8], [9] used autoregressive⁵ code generation to synthesize repairs. Given the sizeable vocabularies in code, many complex dependencies must be accounted for when generating code tokens conditional on previous tokens, the original input tokens, and the compiler error. We simplify the problem into a multi-label classification task here; all that is required is to identify the token positions(s) of the error, and the applicable edit commands. In the vast majority of cases, there is usually only a single change required per line). This allows the model to learn, and rapidly reduce training loss and perform well under test. In addition, the multi-labeling approach (rather than auto-regressive generation also allows us to handle repairs that require multiple fixes on the same line (example below, Fig. 6). It's important to note that a single line can contain several token locations with errors, and distinct edit commands at each position. Limiting the size of the set of possible fixes to 154 will limit the ability to fix identifier names; this is handled by including fix commands that insert and substitute to *unk* in the output vocabulary of LineFix; these fixes are handled by a component is called UNKFix, which is described in §3.5. Note that dealing with multiple fixes on different lines is easily manageable. If there are multiple positions, all with the same fix (like Fig. 6), one can just perform that fix at all the positions. However, for multiple positions and multiple fixes one needs to try all combinations until the javac accepts with no errors. We did not incorporate that to our code, because:

- 1) Trying all possible combinations will slow down the entire process.
 - 2) Two different errors in a line (even in a file) is very rare. In the Blackbox data repository, for example, the majority of files contain just a single syntactical error.

The standard way to train multi-label classification layers is with binary cross-entropy loss (with logits), which is what we use for our fine-tuning. Since both the output layers are closely related to each other, we fine-tuned them simultaneously for 5 epochs. We collected the loss from each layer and added them to define the batch's final loss, and updated the model accordingly. Note that the same pre-trained model parameters (from Fig. 5a) are used to initialize these; during fine-tuning, all parameters in all layers are modified (Fig. 5). We use the Huggingface open-source implementation of RoBERTa [38] for both pre-training and fine-tuning.

Utilizing Compiler Diagnostics during Fine-tuning Apart from localizing the erroneous line, the compiler warning can boost the performance of the fine-tuning model. As an input sequence to the model, we tried two versions, i.e., with the warning, without warning. We observed a small but significant improvement in line-level code fixing (detailed in Section 4.2). Consider the following code snippet from the Blackbox dataset. The variable “bmr” is

5. Autoregressive generation conditions the generation of each token on previously generated tokens, and is used in machine-translation approaches.

```
-System.out.println(*);
+System.out.println("*");
```

Fig. 6. Example requiring two edits to fix.

declared twice, and the second declaration is invalid. Though the javac localizes the error correctly, it is really hard for the model to resolve this without any hint. Our model fails to fix this one when trained without the compiler message. However, with the compiler error message, our RoBERTa-based fine-tuned model can solve errors like this one by deleting the token “double”. This particular example is fixable with a modern IDE; however, it serves as a good illustration of how our model can use error messages. We remind the reader that in general we can handle numerous examples that IDEs cannot. Several typical examples are included in the supplemental file <https://bit.ly/3CMM0TP>.

```
double bmr;
/* some additional irrelevant lines */
boolean isMale = male == 'M';
if(isMale)
    double bmr = ((9.5 * wgt) + (5.0 * hgt)
        + (6.7 * age) + 66.47);
```

Without Warning:

```
double bmr = ((9.5 * wgt) + (5.0 * hgt) + (6.7
* age) + 66.47);
```

With Warning:

```
double bmr = ((9.5 * wgt) + (5.0 * hgt) + (6.7
* age) + 66.47); <SEP> variable declaration
not allowed here
```

LINEFix works best with small sequences. Java is inherently verbose, and so sequence lengths are often beyond the model’s capacity. Compiler diagnostics help us in two ways. Primarily, it helps us localize the error, and second, the message (even if imprecise) helps deep learning models fix the error. This claim is supported by a study (Yasunaga et al. [10]).

3.5 Recovering Unknown Tokens: UNKFix

Recall that LineFix output is restricted to 154 distinct fixes in the fine-tuning model. To deal with edits (inserts or substitutes of identifiers, constants etc.) outside of the limited vocabulary of edits, have an “escape” mechanism. Out of these 154, we included two unique outputs *insert* → *unk* and *substitute* → *unk* to cover other changes. To precisely identify these “unk” tokens, we use UNKFix, which reuses the masked-language model (MLM) we obtained during pre-training. This masked language model can recover the *unk* tokens if sufficient context is given. After getting the position information, we can collect sufficient tokens from the previous and following lines to fill the input buffer, and ask the pre-trained model to unmask the *unk*. Applying this approach, we could fix several *unk*-related program errors like the following ones where the LineFix predicts *insert* → *unk* and *substitute* → *unk* for “Item” and “Integer”, and then the MLM is able to locate them correctly.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on September 23, 2023 at 14:10:12 UTC from IEEE Xplore. Restrictions apply.

```
-public void takeItem (item) {
+public void takeItem (Item item) {
-float number = float.parseInt(text);
+float number = Integer.parseInt(text);
```

Note that though we designed UNKFix primarily for identifiers, it can potentially handle other tokens, including values.

3.6 Integrating SYNShINE into VSCode

To make SYNShINE more broadly accessible, we have made it available within a popular IDE. We have initially chosen VSCode since it’s widely available, free for students⁶, and well-documented; in the future, we will incorporate SYNShINE into other IDEs. The source code for the integration is available in our replication package. A demo video is viewable: <https://youtu.be/AR1nd2PJczU>.

In this VSCode integration, we desired fast response times, and wanted to avoid the requirement for a GPU, since many novices may not have a GPU. So for the SYNShINE deep learning model, we just used CPU floating point operations; to avoid having to reload the (very large) model for each repair request, we wrapped the SYNShINE model within a “correction” server, which services HTTP requests from the IDE.

The IDE triggers a request to SYNShINE when the user requests a fix suggestion. When SYNShINE is triggered, VSCode looks for the active text editor and extracts the (erroneous) code content from there. After getting the content, VSCode sends an HTTP request to the code correction server. Models are pre-loaded in the correction server, so that it can immediately service requests. In this server, the code goes through our proposed pipeline presented in Fig. 2, and the code returns to the editor after finishing all the steps. Now we have two versions of the code, i.e., the buggy code and the corrected version. We highlight the difference and present both versions to the user and allow them to accept or reject the solution.

Note that the demo presented on the link mentioned above was captured on a machine without any GPU. We observe that SYNShINE can operate on a CPU and is quite fast at generating the solution even though the models were trained on GPUs. Just to get a sense of the delay, we randomly chose 200 erroneous programs of various lengths from our dataset, and measured the response time (time from the “SYNShINE” button press to the time the fixed code is received back). The average response time is 0.88 seconds (standard deviation 0.49s, maximum 2.2s). While this by no means instantaneous, we can still provide a fix for a syntax error virtually always within a second or two, potentially saving the novice and instructor’s time. Our approach to integrating SYNShINE into VSCode thus arguably attenuates the need for expensive GPUs, and facilitates the use of the deep learning model in CPU-only machines. The CPU we used for the experiment is “AMD Ryzen 7 2700X”. The code correction server occupies 1.765 GB of the memory.

SynShine’s response time is significantly lower than the time needed by a programmer to fix the program. Brown

6. <https://visualstudio.microsoft.com/students/>

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on September 23, 2023 at 14:10:12 UTC from IEEE Xplore. Restrictions apply.

TABLE 1
Summary Results: Santos et al. Performance is as Reported by Them; we Measured the Others

Santos et al. [8]	DeepFix (short)	DeepFix (long)	SequenceR	BF + FF	SYNShine
46.00%	63.25%	62.14%	56.89%	56.91%	74.89%

and Altadmri divided the mistakes that occurred in the Blackbox repository into 18 different classes, where 11 of them are syntactical errors [39]. The programmers take 13–1000 seconds (median) to fix the mistakes [39]. Our model, on the other hand, takes less than a second on average to process the files and suggest a fix.

4 EVALUATION & RESULTS

In our evaluation, we compare our work with several baselines: Santos et al. DeepFix , BF+FF, and SequenceR. The original DeepFix [7] used a GRU based RNN encoder-decoder translation model, which takes an entire program (with syntax error) as input, and produces a fix. For baselining their BF+FF tool, Ahmed et al. used two versions of DeepFix, one (“short”) trained on error-fix pairs upto 400 tokens long and another (“long”) trained on error-fix pairs upto 800 tokens long. Another approach, SequenceR [40] has reported success in fixing *semantic* errors, when provided with fault localization; it is also adaptable for syntax errors. SequenceR differs from DeepFix in a few ways: it uses a separate fault localizer, and also incorporates a copy mechanism. We describe the intricacies in full detail later. Ahmed et al.’s BF+FF program used a 2-stage transformer-based lenient parser, as described above. Our approach combines several techniques: pre-training, compiler-based reporting, and fine-tuning with novice data.

Below (Table 1), we present summary top-1 accuracy results, evaluated over a random sample of 100,000 examples of length upto 1000 tokens, with single-token errors, taken from the Blackbox dataset. The detailed result is presented in Table 2. We follow the lead of the first paper in the area [8] in this table, reporting performance for single-token

errors, which constitute 57% of the data in Blackbox. We report the numbers for more complex errors below.

As can be seen, SYNShine achieves a substantial performance boost, over all the prior approaches, elevating the performance further and providing us with the motivation to build it into a popular IDE to make it more widely available. Here below, we evaluate the performance in more detail, comparing SYNShine with the closer competitors (we exclude Santos et al. from this comparison) and also examine the contributions of our various stages to the significant overall improvement. We begin with an evaluation of the effect of program length on performance, then we consider the effect of the various components of SYNShine. Finally, we breakdown the performance of SYNShine in repairing various categories of syntax errors.

4.1 Fixing Shorter & Longer Programs

Table 2 baselines the relative performance of SYNShine against prior work, broken down by length, in categories. The rows are different length ranges of programs. The second column is the fraction of the Blackbox programs falling in this length range. The next several columns are baselines from prior work: first two are DeepFix (short) trained on shorter error-fix pairs (upto 400 tokens long), DeepFix (long) trained on pairs up to 800 tokens long. The next two are SequenceR, trained on all pairs in the training set, and BF+FF, trained exactly provided in Ahmed et al.’s scripts. Finally, on the last column we have our results from SYNShine; the 3 columns to the right of the SYNShine column represent the contributions of our 3 components. As can be seen our overall performance exceeds the performance of all the others in every length category, and on the entire sample significantly improves on all of them. Before we examine the numbers in detail, we first present some relevant details on how we measured them.

All evaluations were done on a very large, randomly chosen, representative sample of 100,000 error-fix pairs from Blackbox that were not seen during training by any of the models. The percentages shown in the second column, and the overall performance numbers (all numbers are *top-1 accuracy*) are thus robust estimates of actual performance on programs up to 1000

TABLE 2
Baselining SYNShine Against Prior Work on Syntax Error Correction

Token Range	Percent of Overall Data	DeepFix (short)	DeepFix (long)	SequenceR	BF + FF	SYNShine			Total
						By BLOCKFIX	By LINEFIX	By UNKFIX	
1-100	31.01%	76.71%	73.72%	59.21%	65.16%	21.01%	58.86%	2.41%	82.28%
101-200	29.43%	69.98%	67.15%	57.21%	60.24%	17.53%	58.98%	1.96%	78.47%
201-300	15.25%	63.27%	60.29%	55.40%	54.47%	14.35%	56.00%	1.93%	72.28%
301-400	8.56%	53.71%	54.02%	54.64%	50.01%	10.18%	54.45%	1.89%	66.52%
401-500	5.51%	42.17%	45.47%	54.54%	46.19%	7.71%	54.00%	1.88%	63.59%
501-600	3.63%	32.84%	39.78%	54.47%	42.81%	5.95%	53.83%	2.19%	61.97%
601-700	2.17%	23.76%	33.02%	54.35%	38.07%	3.80%	53.62%	2.10%	59.52%
701-800	1.90%	17.10%	26.57%	53.78%	35.35%	3.04%	51.98%	2.65%	57.67%
801-900	1.34%	11.43%	22.88%	55.56%	32.24%	2.20%	52.84%	2.19%	57.23%
901-1000	1.19%	8.80%	17.94%	53.87%	29.62%	1.27%	51.63%	2.10%	55.00%
Overall		63.25%	62.14%	56.89%	56.91%	15.56%	57.22%	2.11%	74.89%

SequenceR was provided with javac localization.

Authorized licensed use limited to: Univ of Calif Davis. Downloaded on September 23, 2023 at 14:10:12 UTC from IEEE Xplore. Restrictions apply.

tokens long, which constitute around 95% of the Blackbox data. An additional evaluation on a random sample of the entire dataset is reported below. DeepFix (short), DeepFix (long), and BF + FF were all trained and evaluated using the scripts made available in the replication package of Ahmed et al. [9] and Gupta et al. [7].

SequenceR [40] had to be retrained for syntax error correction: Chen et al. originally developed SequenceR for fixing *semantic* bugs, viz., test failures. It uses the OpenNMT translation framework [41] and thus had to be trained using bug-fix pairs. SequenceR assumes that the precise location of the bug was known via fault-localization; the training pairs consisted of a) the buggy region of code, bracketed within `<start_bug> ... <end_bug>` markers, augmented with sufficient context (preceding and succeeding tokens) to make up 1000 tokens of input b) and the corresponding fix, which is the region including the changed code, upto a maximum of 100 tokens; longer fix regions will fail (this almost never happens in our setting). They used an RNN sequence-to-sequence encoder-decoder model that uses LSTM for the recurrent nodes, and incorporates a copy mechanism to enable the model to generate specific local variables, *etc.* in fixes. We used the code provided by Chen et al., and trained the model using Blackbox data; we used the javac compiler to find the error location, and created training/test pairs using the javac indicated location (with context), together with the corresponding novice fix. In our case, since most novices' programs are shorter than 1000 tokens, we provided the entire novice program as context. Once SequenceR is trained, it can generate fixes, given the novice program with error, with location indicated as above. However, SequenceR cannot insert or delete entire lines, so it cannot fix many nesting errors (for example, by inserting or deleting a line with a single "{" or "}" delimiter).

Our overall accuracy ranges between 55% to 82%, and always outperforms DeepFix long (18%-74%), and short (9%-77%), SequenceR (54%-59%) and BF + FF (29%-65%). Both SequenceR and SYNShINE benefit from the error location provided by javac. By improving on prior work at every range, on the entire representative 100,000 sample, SYNShINE achieves significant gains in *overall performance* (bottom line) over the state of the art. Two factors contribute to this improvement: i) javac-based error localization and ii) robustness of LINEFix and UNKFix. javac-based error localization enables a more selective LineFix+UnkFix to the most likely errorful code, thus reducing false positives; Ahmed et al.'s BF + FF attempts corrections throughout the program, resulting in more mistaken corrections. The robustness of LINEFix and UNKFix is really boosted by the pre-training + Fine-tuning strategy; we explore the relative benefits of this step further below.

Table 2, in columns under the SYNShINE header, also shows relative contributions of the components of SYNShINE. First stage is BLOCKFix borrowed from BF + FF. About 20%-25% programs, regardless of length have nesting errors. BLOCKFix's accuracy decreases with program length, and we observe that the contribution of BLOCKFix is low after 700 tokens. However, for the other 75% to 80% programs without nesting errors, LINEFix & UNKFix perform pretty consistently. Finally, we note that 1-1000 tokens cover about 95% of the overall data. To observe the performance of SYNShINE on the overall distribution, including programs over 1000

tokens long, we test it on 5000 random samples. We found that our model can repair 75.36% of the programs, and as before, comfortably exceeds performance of prior tools. Note that if the BLOCKFix model has already fixed the curly braces and there is no other error, javac will not produce any error message, and LINEFix will not process that. Note that we always compare the end-to-end tokens of the reference and the model's proposed sequence; if needless "over" fixes are applied, that will be counted as wrong. Moreover, none of the fixes are credited twice. If the model is fixed by UNKFix, it alone receives credit; we did not count it in the LINEFix column. Likewise, we credited a sample in the LINEFix column, if it is completely fixed by LINEFix and does not receive any help from UNKFix.

We also applied our model on files that required 2 and 3 edits to fix the program and observed 29.4% and 14.4% accuracy, which is much higher than the reported accuracy by Ahmed et al. (19% and 9%). Finally, we note that Ahmed et al. report on a blended strategy where shorter un compilable programs could be sent to DeepFix and longer ones to BF + FF, thus obtaining better performance than either at all lengths. A similar strategy could be employed here, blending SYNShINE with other models, trying all the proposed solutions, and picking the ones that compile. However we didn't implement this approach: we just integrated SYNShINE into VSCode since it performs quite well at all lengths on its own, and avoids the need to load and run many models, and try repeated compiles.

4.2 LINEFix: The Role of Compiler Errors

SYNShINE differs from both versions of DeepFix, and SequenceR, because it's multistage; it differs from BF + FF mainly because of the two new components, LINEFix and UNKFix. We simply reused the BLOCKFix component made available by Ahmed et al.⁷, and find performance very similar to that reported by them for this component. The improvements reported in Table 2 clearly arise from our two new components. We now focus in on LINEFix and evaluate how it contributes to overall performance. LINEFix's task is to take an input line flagged as a relevant syntax error (by javac), together with the actual error, and then output a position, and an editing hint (insert, substitute, delete). LINEFix improves upon the FRAGFix stage of BF + FF in two ways: first, it uses pretraining+finetuning, and second, it also takes the syntax error message from javac as an additional input. The value of pre-training has been extensively documented for code-related tasks [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37], so we focus here on the effect of providing compiler errors. Note again that LINEFix has two tasks: *Localize* the token to be replaced, and *output* an editing command with the correct *Fix*. We evaluate the impact of compiler warnings using 10,000 randomly chosen erroneous lines, of various lengths, each taken with and without the compiler syntax error messages. Since we're evaluating fixing capability on single erroneous lines, rather than entire programs, the numbers reported below are higher than in Table 2.

7. <https://zenodo.org/record/4420845>

TABLE 3
Impact of Using Compiler Error

With compiler error?	Localization	Fix	Complete Correction (Location+Fix)
	F-Score	F-Score	Accuracy
No	90.75%	92.41%	86.71%
Yes	93.58%	93.18%	89.39%

Table 3 presents the impact of using the syntax error message in our tool.

We gain around 2.7% improvement in overall accuracy using the compiler error message. We also see improvements on both Localization and Fix f-scores by providing the compiler message along with the erroneous line (row 1 & 2). The improvement is more for the Localization than for the Fix. We tested the statistical significance of all differences, using Binomial difference of proportions test on a trial sample of 10,000; we then corrected the p-values using Benjamini-Hochberg. The improvements observed when using compiler error message for overall accuracy and fix location f-score are *highly* significant ($p < 1e - 9$); however, the f-score for the fix *per se* are only significantly improved ($0.01 < p < 0.05$). This suggests that the compiler error message is of highly significant help in providing our model with information required to locate the precise token that needs to be edited, and somewhat less so to identify the precise edit that is required. It is *very important* to note however, that the javac compiler is of crucial help in locating the line where the error is located. This above study also shows that the actual error message *per se* helps our model locate the *token* within that line that needs to be edited.

We present an illustrative example of how compiler error messages help. Sometimes the compiler warnings are very precise, e.g., when semicolons or other punctuations are to be inserted. In such cases, it may appear that the task is quite simple, and the model is simply “translating” the error into a fix. We sampled 50 programs and observed how many of them can be fixed just by reading the comments. We observed that in roughly 60% cases, the javac warning is not that helpful, and the model learns to respond in fairly nuanced ways to address the error. Consider the following repair that LINEFix correctly achieves.

```
-return s == reverse ( String s ) ;
+return s == reverse ( s ) ;
```

javac *per se* not helpful: it produces an error message suggesting to insert ")" after "String". LINEFix learns to ignore such messages, and instead correctly omits the token "String". Therefore, the model is not just "translating" the message from javac into a fix; The high capacity of the model, enriched by pre-training and fine-tuning, is deployed to leverage the often incorrect, imprecise message from javac into a good fix. Depending on the error, it can resolve a very imprecise message from javac. Indeed, quite often the same error message from javac can lead the model to provide very different (correct) fixes.

4.3 When SYNShine Fails, and When it Works

We now examine in further detail the cases where SYNShine works correctly, and where it does not. To be conservative, we have defined as a “failure” any fix *not exactly the same* as the one recorded in the Blackbox dataset; note that a) the fix recorded in Blackbox is created by an actual human user, and also b) the recorded fixes always compile without error. We start with an examination of the cases where SYNShine *fails* to produce a correct fix, as per our conservative definition, and then examine in detail the *diversity* of fixes that it does provide.

Fix Failures Despite our over-conservative definition of “failure”, sometimes SYNShine can generate a solution that differs from the user-intended solution but is still compilable with our javac-based compiler. In some cases, the solution is even semantically correct. As an illustration, in Table 6, examples 1, 2 & 3 are fixes generated by SYNShine that not only compile without error, but are also semantically correct. By contrast, the last example in Table 6 is not semantically correct but compilable. Ideally, we’d like to characterize how often SYNShine finds fixes that are not only compilable, but also semantically correct. The *compilability* of a fix that differs from the user’s fix recorded in Blackbox can be determined automatically, and at scale (by

TABLE 4
Compilability of SYNShine

Length	Overall Compilability of fixes	Fixes Exactly Matching Blackbox	Compilability for non-matching cases
1-100	90.18%	82.28%	44.58%
101-200	86.13%	78.47%	35.58%
201-300	79.33%	72.28%	25.43%
301-400	73.35%	66.52%	20.40%
401-500	70.14%	63.59%	17.99%
501-600	67.83%	61.97%	15.41%
601-700	65.92%	59.52%	15.81%
701-800	64.00%	57.67%	14.96%
801-900	63.32%	57.23%	14.24%
901-1000	60.76%	55.00%	13.00%

TABLE 5
Performance of SYNShINE Over Diverse Error Categories

Category	Prevalence of Error Category	Fix Accuracy (in %)
Keyword	5.04%	70.64%
Operator	5.87%	77.73%
Delimiter	80.37%	81.60%
Other	8.72%	60.94%

just compiling!) and we report it below; however, the *semantic correctness* of a fix that differs from a user’s fix requires manual examination, and is not practical to do at a large scale. We try to characterize these to some extent by examining a small sample.

Table 4 presents the overall compilability of the solutions. The second column is the overall *compilability* of the generated fix. This is calculated as the fraction of the number of attempted fixes, that actually results in a successful compilation. The third column is the proportion of fixes that we deem correct, based on *exact* match with the fix recorded in Blackbox (the numbers will match shown in the right-most column of Table 2). As can be seen, we record many compilable cases as incorrect. The last column in Table 4 shows the proportion of apparent failures that are actually compilable: as an illustration, for programs up to 100 tokens long, about 45% of the cases that we record as an incorrect fix, in fact compile correctly. Depending on length, between 13% and 45% of the fixes we classify as failures are actually compilable. Table 6, examples 1,2,3,4 are exactly such fixes.

Now what proportion of these “compilable failures” are actually semantically correct? To get a (very) rough estimate of this, we did a small manual study. We randomly collect 50 cases where the model generates a compilable fix, that fails to match the user fix recorded in Blackbox. We found that about 18% of programs are semantically correct.

To summarize: even in our very conservative evaluation, SYNShINE produces the same fixes as recorded by a human in a sizable fraction (roughly 75%) of errors in our novice dataset; an examination of SYNShINE’s failures suggests that it could possibly be helpful in some additional cases.

Fix Diversity What kinds of errors does SYNShINE fix? In our dataset, about 80% of the errors are related to delimiters, and even solving only those would make a significant dent. However, the novices make syntax errors in using keywords, operators, identifiers, and numbers; sometimes they introduce illegal spaces, declarations, characters, etc. We examined how SYNShINE performs with respect to different types of errors. For convenience, we divided the error into four major categories- keywords (all Java keywords), delimiters (e.g., semicolon, comma, parentheses, braces, brackets), operators (all Java operators), and others (identifiers, literals, and anything that falls outside the first three categories).

To do categorization, we followed two rules. Errors that required substitutes or inserts belonged to the category of the substituted or inserted token; errors that required deletion belonged to the category of the deleted token. Thus if an error required a semicolon to be inserted, it was in the “delimiter” category; if an error required an extra “if” keyword to be deleted, it was in the “keyword” category.

We randomly sampled a 5K test dataset, and determined the error category prevalence in this dataset; see Table 5, first column, for the prevalence of errors in various categories. Delimiter errors dominate, and thus our model learns to fix those best (81.6% accuracy); however, it performs well in other categories (60%-78% accuracy). The take-away from this analysis is that SYNShINE performs reasonably well at a wide range of syntax errors.

5 RELATED WORK

The most closely related works are DeepFix [7], BF + FF [9], and Santos et al. [8] which we have discussed above. We also discussed SequenceR [40]. We have compared SYNShINE to all of these.

Gupta et al. [12] applied reinforcement learning to a very similar dataset like DeepFix [7]. It utilizes total count of compiler errors as a part of the reward mechanism. However, RLAssist [12] shows only a very minor improvement over DeepFix [7], and also it takes the whole program as input. Therefore, we did not re-implement RLAssist [12]. Though RLAssist [12] looks into compiler errors but it does not directly uses the error messages as we do. DeepDelta [42] is another approach that fixes compiler errors but mostly identifier name-related errors, not syntax errors. DeepDelta [42] was developed and tested on code from professional developers at Google. The authors also assume that precise knowledge of the location will be given to the program. Yasunaga et al. [10], [13] introduce two compiler-dependent approaches to fix C program: DrRepair that utilizes C compiler warnings with a graph-based self-supervised approach, and BIFI that applies two models “critic” and “fixer” to fix the programs. A tool for the C programming language, Tracer, abstracts the code and uses a seq2-seq model on the source code abstractions that are later concretized [43].

All the DNN based Automatic Program Repair (APR) tasks have a fault localization step [40], [44], [45], [46], [47], and these tools’ performance depends a lot on the fault-localizer. Semantic code correction is an inherently difficult problem, and syntax correction can be considered as a subset of semantic code correction problems. None of the previous syntax correction tools has compared their work with these tools because previous syntax correction tools did not depend on any fault localizer. Some of the APR tools [45],

TABLE 6
Examples Showing the Compilability of the Model

Seq No	Buggy Line	Model	Original Fixed
1	int i = ((int) (Math . random () * 3);	int i = ((int) (Math . random () * 3));	int i = (int) (Math . random () * 3);
2	int userInt_1 , int userInt_2 ;	int userInt_1 ; int userInt_2 ;	int userInt_1 , userInt_2 ;
3	System . out . print (" Hello, world. ");	System . out . print (" Hello, world. ");	System . out . print (" Hello, world. ");
4	System . out . println (" sum = " + (sum +));	System . out . println (" sum = " + (sum));	System . out . println (" sum = " + (sum + 5));

[48], [49], [50] expects syntactically correct programs and those approaches are not applicable for syntactical code correction. For our purposes the most directly compatible recent APR tool was “SequenceR” [40] which reported good performance, and also fixes errors at the line level; it was readily adapted to using the javac to locate the line to be fixed, so we chose it for comparison. Pradel et al. also detect specific types of bugs (e.g., accidentally swapped function arguments, incorrect binary operators, and incorrect operands in binary operations) but in syntactically correct code [51].

Brown et al. used BlueJ IDE to collect the data in Blackbox repository [11]. In this paper, we did a case study on the performance of the popular IDEs (e.g., Eclipse, IntelliJ, VSCode, BlueJ) in fixing novice programs. We compare repair hints from Eclipse JDT Core Compiler for Java (ECJ) (used in both Eclipse and VSCode) and javac (used by IntelliJ and BlueJ). That is, both Eclipse and VSCode present the same error messages, and IntelliJ and BlueJ present the same error messages. Four IDEs, but ultimately, only two compilers. SYNShine improves upon repair hints from both compilers. Therefore, we primarily focus on Eclipse and IntelliJ for the case study. We chose VSCode because it is popular, well-documented, available free for students, and is easy to extend. We were able to integrate SYNShine into VSCode without any major difficulties.

6 CONCLUSION

We have described SYNShine, a machine-learning based tool to fix syntax errors in programs. SYNShine leverages RoBERTa pre-training, uses compiler errors (both location and message), and generates fixes using multi-label classification, rather than autoregressive generation, to achieve substantial improvements in fixing syntax errors. Our evaluation shows substantial improvements in fixing rates over the previous best results reported by BF+FF, and other tools, at all program lengths. Our evaluations suggest that the use of compilers to locate the precise line provides a big advantage; our evaluations also suggest that the compiler error message *per se* may be helpful in locating the precise token within the line that needs to be repaired. We have built SYNShine into the VSCode IDE, and have found that even without a GPU, the SYNShine-enhanced VSCode can fix syntax errors fairly quickly, often in less than a second. We have made all the source-code and data available, to the extent allowable under UK Law applicable to the Blackbox dataset. SYNShine can fix errors that IDEs (Eclipse, IntelliJ, BlueJ, and VSCode) cannot. In the supplementary materials (<https://bit.ly/3CMM0TP>) we show several real-world examples of novice-made errors that cannot be fixed by any of these IDEs, but can be fixed by SYNShine. Finally, the entire source for our SYNShine, including the VSCode extension, is made available anonymously at <https://doi.org/10.5281/zenodo.4563241>.

ACKNOWLEDGMENTS

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Ahmed was also supported by UC Davis College of Engineering Dean’s Distinguished Fellowship.

REFERENCES

- [1] J. C. Spohrer and E. Soloway, “Novice mistakes: Are the folk wissdoms correct?,” *Commun. ACM*, vol. 29, no. 7, pp. 624–632, 1986.
- [2] M. McCracken et al., “A multi-national, multi-institutional study of assessment of programming skills of first-year cs students,” in *Proc. Work. Group Rep. ITiCSE Innov. Technol. Comput. Sci. Educ.*, 2001, pp. 125–180.
- [3] E. Lahtinen, K. Ala-Mutka, and H.-M. Järvinen, “A study of the difficulties of novice programmers,” *ACM SIGCSE Bull.*, vol. 37, no. 3, pp. 14–18, 2005.
- [4] J. Jackson, M. Cobb, and C. Carver, “Identifying top java errors for novice programmers,” in *Proc. IEEE Front. Educ. 35th Annu. Conf.*, 2005, pp. T4C–T4C.
- [5] S. K. Kümmerfeld and J. Kay, “The neglected battle fields of syntax errors,” in *Proc. 5th Australas. Conf. Comput. Educ.*, 2003, pp. 105–111.
- [6] P. Denny, A. Luxton-Reilly, and E. Tempero, “All syntax errors are not equal,” in *Proc. 17th ACM Annu. Conf. Innov. Technol. Comput. Sci. Educ.*, 2012, pp. 75–80.
- [7] R. Gupta, S. Pal, A. Kanade, and S. Shevade, “DeepFix: Fixing common c language errors by deep learning,” in *31st AAAI Conf. Artif. Intell.*, 2017, pp. 1345–1351.
- [8] E. A. Santos, H. V. Campbell, D. Patel, A. Hindle, and J. N. Amaral, “Syntax and sensibility: Using language models to detect and correct syntax errors,” in *Proc. IEEE 25th Int. Conf. Softw. Anal. Evol. Reengineering*, 2018, pp. 311–322.
- [9] T. Ahmed, P. Devanbu, and V. Hellendoorn, “Learning lenient parsing & typing via indirect supervision,” *Empirical Softw. Eng.*, vol. 26, 2021, pp. 1–31.
- [10] M. Yasunaga and P. Liang, “Graph-based, self-supervised program repair from diagnostic feedback,” in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 10 799–10 808.
- [11] N. C. C. Brown, M. Kölling, D. McCall, and I. Utting, “Blackbox: A large scale repository of novice programmers’ activity,” in *Proc. 45th ACM Tech. Symp. Comput. Sci. Educ.*, 2014, pp. 223–228.
- [12] R. Gupta, A. Kanade, and S. Shevade, “Deep reinforcement learning for programming language correction,” 2018, *arXiv: 1801.10467*.
- [13] M. Yasunaga and P. Liang, “Break-it-Fix-it: Unsupervised learning for program repair,” 2021, *arXiv:2106.06600*.
- [14] R. P. Medeiros, G. L. Ramalho, and T. P. Falcão, “A systematic literature review on teaching and learning introductory programming in higher education,” *IEEE Trans. Educ.*, vol. 62, no. 2, pp. 77–90, May 2018.
- [15] C. Watson and F. W. Li, “Failure rates in introductory programming revisited,” in *Proc. Conf. Innov. Technol. Comput. Sci. Educ.*, 2014, pp. 39–44.
- [16] J. Bennedsen and M. E. Caspersen, “Failure rates in introductory programming,” *AcM SIGcSE Bull.*, vol. 39, no. 2, pp. 32–36, 2007.
- [17] B. A. Becker et al., “Compiler error messages considered unhelpful: The landscape of text-based programming error message research,” in *Proc. Work. Group Rep. Innov. Technol. Comput. Sci. Educ.*, 2019, pp. 177–210.
- [18] T. Barik et al., “Do developers read compiler error messages?,” in *Proc. IEEE/ACM 39th Int. Conf. Softw. Eng.*, 2017, pp. 575–585.
- [19] E. Kantorowitz and H. Laor, “Automatic generation of useful syntax error messages,” *Softw. Pract. Experience*, vol. 16, no. 7, pp. 627–640, 1986.
- [20] T. Schorsch, “Cap: An automated self-assessment tool to check pascal programs for syntax, logic and style errors,” *ACM SIGCSE Bull.*, vol. 27, no. 1, pp. 168–172, 1995.
- [21] M.-T. Luong, H. Pham, and C. D. Manning, “Effective approaches to attention-based neural machine translation,” 2015, *arXiv:1508.04025*.
- [22] A. Vaswani et al., “Attention is all you need,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5998–6008.
- [23] R.-M. Karapatsis, H. Babii, R. Robbes, C. Sutton, and A. Janes, “Big code != big vocabulary: Open-vocabulary models for source code,” in *Proc. IEEE Int. Conf. Softw. Eng.*, 2020, pp. 1073–1085.
- [24] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” 2018, *arXiv:1810.04805*.
- [25] Y. Liu et al., “RoBERTa: A robustly optimized bert pretraining approach,” 2019, *arXiv:1907.11692*.
- [26] Z. Feng et al., “CodeBERT: A pre-trained model for programming and natural languages,” 2020, *arXiv:2002.08155*.
- [27] A. Kanade, P. Maniatis, G. Balakrishnan, and K. Shi, “Learning and evaluating contextual embedding of source code,” in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 5110–5121.

- [28] E. Biswas, M. E. Karabulut, L. Pollock, and K. Vijay-Shanker, "Achieving reliable sentiment analysis in the software engineering domain using bert," in *Proc. IEEE Int. Conf. Softw. Maintenance Evol.*, 2020, pp. 162–173.
- [29] T. Zhang, B. Xu, F. Thung, S. A. Haryono, D. Lo, and L. Jiang, "Sentiment analysis for software engineering: How far can pre-trained transformer models go?," in *Proc. IEEE Int. Conf. Softw. Maintenance Evol.*, 2020, pp. 70–80.
- [30] S. Lu et al., "CodeXGLUE: A machine learning benchmark dataset for code understanding and generation," 2021, *arXiv:2102.04664*.
- [31] B. Roziere, M.-A. Lachaux, M. Szafraniec, and G. Lample, "DOBF: A deobfuscation pre-training objective for programming languages," 2021, *arXiv:2102.07492*.
- [32] A. Mastropalo et al., "Studying the usage of text-to-text transfer transformer to support code-related tasks," 2021, *arXiv:2102.02017*.
- [33] K. Jesse, P. T. Devanbu, and T. Ahmed, "Learning type annotation: Is big data enough?," in *Proc. 29th ACM Joint Meeting Eur. Softw. Eng. Conf. Symp. Found. Softw. Eng.*, 2021, pp. 1483–1486.
- [34] T. Ahmed, P. Devanbu, and A. A. Sawant, "Learning to find usage of library functions in optimized binaries," *IEEE Trans. Softw. Eng.*, to be published, doi: [10.1109/TSE.2021.3106572](https://doi.org/10.1109/TSE.2021.3106572).
- [35] D. Guo et al., "GraphCodeBERT: Pre-training code representations with data flow," in *Proc. Int. Conf. Learn. Representations*, 2020, pp. 1–18.
- [36] W. Qi et al., "ProphetNet-X: Large-scale pre-training models for english, chinese, multi-lingual, dialog, and code generation," 2021, *arXiv:2104.08006*.
- [37] W. Ahmad, S. Chakraborty, B. Ray, and K.-W. Chang, "Unified pre-training for program understanding and generation," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, 2021, pp. 2655–2668.
- [38] Huggingface, "Huggingface transformers," 2018. [Online]. Available: <https://github.com/huggingface/transformers>
- [39] N. C. Brown and A. Altadmir, "Novice java programming mistakes: Large-scale data versus. educator beliefs," *ACM Trans. Comput. Educ.*, vol. 17, no. 2, 2017, Art. no. 7.
- [40] Z. Chen, S. J. Kommrusch, M. Tufano, L.-N. Pouchet, D. Poshyvanyk, and M. Monperrus, "SequenceR: Sequence-to-sequence learning for end-to-end program repair," *IEEE Trans. Softw. Eng.*, vol. 47, no. 9, pp. 1943–1959, Sep. 2021.
- [41] G. Klein, Y. Kim, Y. Deng, J. Senellart, and A. M. Rush, "OpenNMT: Open-source toolkit for neural machine translation," 2017, *arXiv:1701.02810*.
- [42] A. Mesbah, A. Rice, E. Johnston, N. Glorioso, and E. Aftandilian, "Deepdelta: Learning to repair compilation errors," in *Proc. 27th ACM Joint Meeting Eur. Softw. Eng. Conf. Sympos. Foundations Softw. Eng.*, 2019.
- [43] U. Z. Ahmed, P. Kumar, A. Karkare, P. Kar, and S. Gulwani, "Compilation error repair: For the student programs, from the student programs," in *Proc. 40th Int. Conf. Softw. Eng. Softw. Eng. Educ. Training*, 2018, pp. 78–87.
- [44] M. Tufano, J. Pantiuchina, C. Watson, G. Bavota, and D. Poshyvanyk, "On learning meaningful code changes via neural machine translation," in *Proc. IEEE 41st Int. Conf. Softw. Eng.*, 2019, pp. 25–36.
- [45] Y. Li, S. Wang, and T. N. Nguyen, "DLFix: Context-based code transformation learning for automated program repair," in *Proc. 42th Int. Conf. Softw. Eng.*, 2020, pp. 602–614.
- [46] Y. Ding, B. Ray, P. Devanbu, and V. J. Hellendoorn, "Patching as translation: The data and the metaphor," in *Proc. IEEE/ACM 35th Int. Conf. Automated Softw. Eng.*, 2020, pp. 275–286.
- [47] T. Lutellier, H. V. Pham, L. Pang, Y. Li, M. Wei, and L. Tan, "Coconut: Combining context-aware neural translation models using ensemble for program repair," in *Proc. 29th ACM SIGSOFT Int. Symp. Softw. Testing Anal.*, 2020, pp. 101–114.
- [48] C. Le Goues, T. Nguyen, S. Forrest, and W. Weimer, "GenProg: A generic method for automatic software repair," *IEEE Trans. Softw. Eng.*, vol. 38, no. 1, pp. 54–72, Jan./Feb. 2011.
- [49] S. Chakraborty, M. Allamanis, and B. Ray, "Tree2Tree neural translation model for learning source code changes," 2018, *arXiv:1810.00314*.
- [50] F. Long and M. Rinard, "Automatic patch generation by learning correct code," in *Proc. 43rd Annu. ACM SIGPLAN-SIGACT Symp. Princ. Program. Lang.*, 2016, pp. 298–312.
- [51] M. Pradel and K. Sen, "DeepBugs: A learning approach to name-based bug detection," *Proc. ACM Program. Lang.*, vol. 2, no. OOPSLA, pp. 1–25, 2018.



Toufique Ahmed received the BSc and MSc degrees in computer science and engineering from the Bangladesh University of Engineering and Technology (BUET), in 2014 and 2016, respectively. He is currently working toward the PhD degree with University of California, Davis (UC Davis). His research interests include Software Engineering, the Naturalness of Software, Machine Learning, and Sentiment Analysis. He is the recipient of the five-year prestigious Dean's Distinguished Graduate Fellowship (DDGF) offered by The Office of graduate studies, The College of Engineering, and the Graduate Group in Computer Science, UC Davis.



Noah Rose Ledesma is currently working toward the undergraduate degree in computer science student with the UC Davis, and a research assistant with the DECAL lab. He is interested in machine learning, computer graphics, autonomous vehicles, and other subjects. Recently, Noah participated in an internship in which he developed infrastructure for Google Cloud's artificial intelligence platform. He will begin his career as a full-time software engineer after he graduates, in June 2022.



Premkumar Devanbu received the BTech degree from IIT Madras and the PhD degree from Rutgers University. He is currently distinguished professor of computer science with UC Davis. His research interests include Empirical Software Engineering and the applications of the Naturalness of Software. He is an ACM Fellow.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.