TraceFixer: Execution Trace-Guided Program Repair

Islem Bouzenia
fi_bouzenia@esi.dz
University of Stuttgart
Germany

Yangruibo Ding
yrbding@cs.columbia.edu
Columbia University
USA

Kexin Pei
kpei@cs.columbia.edu
Columbia University
USA

Baishakhi Ray
rayb@cs.columbia.edu
Columbia University
USA

Michael Pradel
michael@binaervarianz.de
University of Stuttgart
Germany

ABSTRACT

When debugging unintended program behavior, developers can often identify the point in the execution where the actual behavior diverges from the desired behavior. For example, a variable may get assigned a wrong value, which then negatively influences the remaining computation. Once a developer identifies such a divergence, how to fix the code so that it provides the desired behavior? This paper presents TraceFixer, a technique for predicting how to edit source code so that it does not diverge from the expected behavior anymore. The key idea is to train a neural program repair model that not only learns from source code edits but also exploits excerpts of runtime traces. The input to the model is a partial execution trace of the incorrect code, which can be obtained automatically through code instrumentation, and the correct state that the program should reach at the divergence point, which the user provides, e.g., in an interactive debugger. Our approach fundamentally differs from current program repair techniques, which share a similar goal but exploit neither execution traces nor information about the desired program state. We evaluate TraceFixer on single-line mistakes in Python code. After training the model on hundreds of thousands of code edits created by a neural model that mimics real-world bugs, we find that exploiting execution traces improves the bug-fixing ability by 13% to 20% (depending on the dataset, within the top-10 predictions) compared to a baseline that learns from source code edits only. Applying TraceFixer to 20 real-world Python bugs shows that the approach successfully fixes 10 of them.

1 INTRODUCTION

When trying to fix a bug, a developer must localize the problem and then edit the source code to prevent the problem from happening. During the first step, a developer typically identifies a source code location and understands how the behavior at this location differs from the desired behavior. For example, when stepping through a program in an interactive debugger, a developer may realize a point in the execution where the value of some variable differs from the value the variable should have. We refer to the point during an execution where the actual behavior diverges from the desired behavior as the *divergence point* of a bug, and we call the state the program should reach at this point the *desired state*. Given these two pieces of information, the second step is to edit the source code in a way that prevents divergence and reaches the desired state.

For example, Figure 1 shows a snippet of code taken from the deepmind/pysc2 project[1], where the program behavior diverges at the last shown line. The middle part of the figure gives the trace obtained by executing the buggy code, while the right part shows the trace the developer intends the program to produce. The difference is in the state produced at the divergence point: Instead of storing `[88, 89, 90, 91, 92]` in variable `ports`, the desired state is that `ports` has the value `[88, 89, 90, 91]`.

Automated program repair has made impressive progress in fixing programming mistakes over the past few years [31], and learning-based approaches are shown to be the current state of the art [10, 15, 36, 69]. However, none of these approaches considers runtime behavior, such as a partial execution trace or the desired state, as an input to the approach. Instead, many repair approaches validate candidate repairs against test cases of the expected behavior as a post-processing step after predicting fix candidates [10, 30, 36, 60, 68] or use tests as feedback for repair predictions [65]. As illustrated by the example in Figure 1, determining the intended fix without knowing what state the program should reach is clearly difficult. Providing such information to a learning-based repair technique hence offers an opportunity to further improve their effectiveness.

This paper presents TraceFixer, the first learning-based program repair technique guided by execution traces. By analogy to the human way of debugging, our approach considers a program's execution trace alongside its code. The inputs to the approach are the source code of a buggy program, an excerpt of the execution trace produced by this program, and the desired state at the divergence point. Given these inputs, the approach predicts a variant of the code that produces the desired trace. TraceFixer is enabled by a neural model that reasons about both code and corresponding execution traces to predict the proper fix.

In the example of Figure 1, TraceFixer starts by instrumenting and executing the program using, e.g., inputs provided by a user or a test case. The result is an execution trace, as shown in the middle part of the figure. Based on this execution, the developer determines the divergence point and gives the desired state of the program at this point. For the example, the developer indicates that the last shown line is the divergence point, and that the `ports` variable should have the value `[88, 89, 90, 91]`. Given the incorrect code, its execution trace, and the desired state, TraceFixer predicts how to fix the code by suggesting to replace the last shown line

---

[1]https://github.com/deepmind/pysc2/commit/803de8a0f707efc3967fde683836c202c73215af

```
Buggy code                                          Trace                              Desired trace

# some code before
while True:
    start_port = portpicker.pick_unused_port()      ENTER_WHILE while True:             ENTER_WHILE while True:
                                                    start_port = 88                    start_port = 88
    ports = [start_port + p for p in range(5)]      ports = [88, 89, 90, 91, 92]       ports = [88, 89, 90, 91]
```

Figure 1: Example of a buggy program with corresponding execution traces.

with `ports = [start_port + p for p in range(4)]`. Indeed, with this fix, the program will produce the desired state.

Creating a trace-guided neural repair approach leads to two key challenges not addressed in prior work. (C1) First, since our approach is data-driven, we need an extensive dataset that consists of pairs of buggy code and the corresponding correct code, along with execution traces of their behavior. (C2) Second, we need to identify a suitable neural model and an effective format for feeding the different input modalities–buggy code, execution trace, and desired state–into the model.

TraceFixer addresses challenge C1 through a neural bug injection model that creates hundreds of thousands of realistic bugs from code with existing test cases. Inspired by the success of existing neural approaches to generate bugs [4, 39, 41], we train a neural bug injection model for Python that mimics single-line bugs observed to occur in the wild [27]. To address challenge C2, TraceFixer provides an input formatting module that represents the different modalities in a way that is compatible with popular sequence-to-sequence models, allowing us to build upon a strong pre-trained model [58].

Our evaluation applies TraceFixer to single-line bugs in Python. We inject hundreds of thousands of realistic bugs into two large-scale datasets of executable programs and then train the repair model to fix these bugs. Once trained, the model successfully repairs 57%–63% of the bugs with its top-most suggestion, and 82%–87% of the bugs within the top-10 suggestions. Out of a set of 20 real-world Python bugs, TraceFixer successfully fixes 10. Finally, we compare our approach with a neural baseline model that, similar to prior techniques [10, 15, 36, 69], learns only from code but not runtime information, and find that TraceFixer improves upon the state of the art by 13% to 20%.

While practical techniques for finding the divergence point and obtaining the traces provided to TraceFixer are beyond the scope of this paper, we envision at least two ways of using our approach in practice. First, we believe that appropriate tool support will make it relatively easy for developers to provide this information when they are anyway localizing a bug. For example, an interactive debugger that allows developers to step through their code and inspect runtime values could provide an option for marking a specific value as wrong and for providing the correct value instead. Second, the divergence point and traces could be gathered automatically with existing fault localization techniques [62] applied to test cases that expose the problem, which generate-and-validate repair techniques [31] commonly assume to exist.

In summary, this paper contributes the following:

- An automated repair technique that learns not only from source code but also from runtime traces and information about the state a program should reach.

- A neural repair model enabled by learning from a large-scale set of automatically injected bugs and a novel input formatting that feeds different modalities into a sequence-to-sequence architecture.
- Empirical evidence that considering execution traces improves the bug-fixing abilities of the approach and that Trace-Fixer is effective for real-world Python bugs.

2 APPROACH

In this section, we present our data-driven, deep learning-based automated program repair approach called TraceFixer. We start by defining the problem we address (Section 2.1), then give an overview of the approach (Section 2.2), and finally explain its different components in detail (Sections 2.4 to 2.7). We also give a preliminary background in Section 2.3

2.1 Problem Definition

TraceFixer reasons about executions of programs, which we represent as a sequence of states:

Definition 2.1 (State). The state s at a point during the execution of a program p is a list of pairs (n, v) where the n is the name of a variable available in the current scope during the execution of p and v is the value that n refers to.

We consider both primitive values, e.g., integers and strings, and non-primitive values, e.g., a list or an instance of a class. The approach focuses on the state in the program's main memory and ignores any other state, such as the state of the file system, the state reachable via the network, or the state of another process.

During execution, a program transfers from one state to another. We represent these transfers in an execution trace:

Definition 2.2 (Execution trace). The execution trace of program p is a pair (s_{init}, E), where s_{init} is the initial state and E is a sequence of events observed starting from the initial state. Each event in E is a pair (l_i, s_i) where l_i is an executed line of code in p, and s_i is the state after executing the line.

The full execution trace until triggering a bug may consist of many more events than a developer would typically inspect when reasoning about an incorrect execution. Instead of considering the full execution trace starting at the first executed line of the program, the initial state in a trace may be any point in the program's execution. In our evaluation, we consider traces with up to three events. If the line l_i captured in an execution trace involves a function call, then we include the return value of the call (if any) as part of the state s_i.
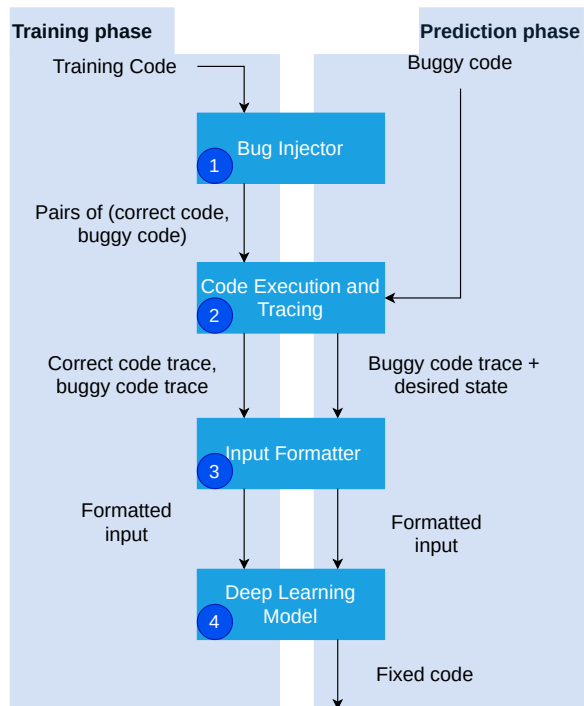
Figure 2: Overview of TraceFixer approach.

The usage scenario of TraceFixer is that a developer identifies the point during an execution where the actual and the intended execution start to diverge:

Definition 2.3 (Divergence point). Consider two traces (s_{init}, E) and (s'_{init}, E') with $E = [(l_1, s_1), ..., (l_k, s_k)]$, $E' = [(l'_1, s'_1), ..., (l'_k, s'_k)]$, $l_i = l'_i$ and $s_i = s'_i$ for all $i < k$, but $l_k \neq l'_k$ or $s_k \neq s'_k$. That is, the traces are equal except for the last event, where either the line that gets executed, what state is reached after executing the line, or possibly both differ. The last event in the two traces is called the divergence point.

Once a developer has identified a divergence point, e.g., while stepping through a program using an interactive debugger, the problem addressed in this paper is how to fix the incorrect code:

Definition 2.4 (Trace-guided repair problem). Given a buggy piece of code c, a trace t produced by this code, and a state s that is the desired state of the program at the divergence point, the problem is to predict the fixed code c' that, if executed, will yield a trace that differs from t only by reaching state s.

2.2 Overview

Our approach to addressing the abovementioned problem consists of four main steps represented in Figure 2. (i) *Bug Injector.* The first step is specific to the training phase. In every deep learning-based application, having a training dataset is crucial. In our case, we first need a dataset composed of pairs of buggy code and the corresponding correct code. Therefore, we design a neural bug injector: given a correct code and a location, this step injects bugs by

modifying the code at the given location. We leverage a pre-trained seq2seq model, which receives a correct program with the injection location (target line) as input and outputs the buggy version of the target line. At a high level, this step can be thought of using a seq2seq model in the reverse manner of neural bug fixing [8]. The injector further replaces the target line in the correct program with the generated buggy line. (ii) *Code Execution and Tracing.* Once we have pairs of correct and buggy code, we instrument and execute them to get their execution traces. In the training phase, the approach extracts the divergence point and the desired state from these traces. In the prediction phase, the developer directly provides this information. (iii) *Input Formatter.* The buggy code, the trace of the buggy code, and the desired state will be passed into a neural model. To this end, we format the different data modalities in a way suitable for seq2seq architectures. (iv) *Deep Learning Model.* In the final step, we train a neural model to fix the buggy code and, once trained, use the model to predict the fix for previously unseen bugs.

2.3 Preliminary

Transformer-based sequence-to-sequence model. In recent years, Transformer based sequence-to-sequence (seq2seq) models have been applied to code editing tasks [1, 8, 15, 58, 66]. The model contains an encoder and a decoder, and both are built based on a stack of Transformer layers [55]. The encoder summarizes the information of the input sequence, and a decoder generates the new token sequences based on the output of the encoder. The decoder predicts the new sequence left-to-right, token by token. Every token is generated conditioned on both the encoded input and the generated prefix (i.e., left tokens already generated in the previous time steps). The model's input is the flattened sequence of a program, and the expected output is the edited code so that the model will learn how to change the input to match the ground truth during training.

CodeT5. CodeT5 [58] is a state-of-the-art pre-trained seq2seq model for source code. The pretraining is done across different code-generation tasks, including code summarization, code generation, and code refinement. The code refinement task is the most related task to our work: it randomly masks code tokens of the input with a special token <MASK>, and trains the model to predict the masked tokens back. Conceptually, this task teaches the model to edit the code, which can potentially assist with our proposed models to inject or fix bugs. CodeT5 is pre-trained with millions of samples with more than 100 epochs. With such an effective pretraining, it reports promising results in both code understanding and generation tasks. In addition, CodeT5 provides two model sizes in the original paper: small and base. The former has six layers of transformers for both encoder and decoder, while the latter has 12 layers for both encoder and decoder.

Byte-pair Encoding. The vocabulary for source code is essentially open-ended since developers could name identifiers with any preferred combination of words or introduce new words that are not part of any natural language's vocabulary [28]. To train a neural network on source code, however, we need a fixed length of vocabulary [23] while minimizing the out-of-vocabulary cases. Byte-pair

encoding (BPE) was initially applied to address the rare words issue in neural machine translation [51]: the algorithm splits the rare and unknown words into sub-words already part of the vocabulary; in the worst case, a complicated word will be split into characters. Similarly, we apply BPE to address the out-of-vocabulary issue of source code in our models: the common code tokens, such as keywords, will be encoded as single tokens, while the complicated code tokens will be split into common sub-tokens in the vocabulary.

2.4 Bug Injection

To train a neural model to repair bugs, first, we need a dataset composed of the buggy and correct code pairs. The dataset should be large enough to enable effective learning by a neural model. Each buggy and correct code pair in the dataset should be executable to allow gathering traces. Last but not least, the injected bugs should resemble real-world bugs. To this end, we use a neural bug injection technique, motivated by the success of existing neural approaches to generate realistic bugs [4, 39, 41].

The backbone of our tool is a Transformer-based, sequence-to-sequence (seq2seq) neural network. To save the training efforts and ensure the code generation quality, we initialize our neural network with CodeT5-small [58], which has been trained on millions of code samples with a code refinement task, so it already has good knowledge about how to edit code. Consequently, loading it as the initial model can help our neural bug injector converge faster during training and generate more valid programs than random initialization.

We use the PySStuBs dataset [27] to train the neural bug injector. This dataset contains samples with both buggy and patched versions, and we build the model input with the patched version while the ground truth is the corresponding bug. Consequently, we train the neural model to inject bugs by reversing the bug-fixing process. Furthermore, we ignore the samples that cannot be checked out due to the invalid commit or file hashes, and we do not consider the files written with Python2. We end up with 47,314 valid samples; we split them into 95% and 5% portions for training and tuning parameters, respectively.

The workflow of the bug injector is shown in Figure 3. The first step of the tool is to create the input for the neural model based on a correct program and a target line at which to inject bugs. The correct program will be a patch in the PySStuBs dataset, and the target line is the changed line when fixing the corresponding bug. The bottom-left of Figure 3 shows an example of the model input, where `log.log_open(logFile, startTime)` is the target line. To ensure the model performance, we design the model input following the state-of-the-art neural code editor, MODIT[8], where we extract the line to be edited from the program and prepend it at the beginning of the input sequence. We use special tokens to split the target line and its contexts so the model will know their positions. To further indicate the original location of the target line and align with the design of the pretrained CodeT5, we fill that location with a placeholder token, `<MASK>`.

Given the model input, the second step is to train the seq2seq model to predict the buggy version of the target line. During training, the model will be fed with a ground-truth bug, and the model learns to construct the bug left-to-right, token-by-token. Later, we

replace the `<MASK>` in the input with the model prediction to generate the buggy program.

2.5 Execution and Tracing

Since our approach is based on program repair using traces, getting those traces is essential and required in both the training and the prediction phase. For that, we instrument the target code and then execute it. Our instrumentation targets the line level, meaning that we collect the program state after executing every line of code which matches definition 2.2. However, we limit the instrumentation and the collected trace to a specific scope. In our case, the scope is usually one function or a sequential script consisting of one file. We intentionally do not trace subsequent calls to other functions or modules, mimicking a developer who steps through the execution of a function without stepping into callees. In the case of a return statement, we copy the value of the returned variable before it is returned, and we copy the new state if the return expression changes it.

When collecting the traces, we save the printable string of all variables in scope in a format similar to the one of definition 2.1. For example, if at some point of execution we have `primes = [2, 3, 5, 7, 11]`, the saved state related to the variable `primes` would be `[("primes", "[2, 3, 5, 7, 11]")]`. Saving the evaluation of the variable "primes" as a string is easy because the printable string of the variable primes matches the value that a developer cares about, which is, in this case, the list of numbers themselves. Unfortunately, not all variables are printable like primitive types and standard data structures. For that, if the value of a variable is not a primitive type or a standard data structure, we inspect the object referred to by the variable. We collect the values of all printable attributes of that object. For example, after this line of code: `rectangle = Rectangle(2, 5)`, the variable `rectangle` will refer to an instance of the class `Rectangle`, which has two attributes: `width` and `height`. The state of the program after that line would be `[("rectangle.width", "2"), ("rectangle.height", "5")]`. In the case where there are no printable values when inspecting the attributes of an object, we show the type of the variable. For example, if the variable `rectangle` has no printable attributes then the trace of the line `rectangle=Rectangle(2, 5)` would be `[("rectangle", "instance(Rectangle)")]`.

In the following, we describe the flow of instrumentation and execution to get the traces.

Training phase. We need the traces of both buggy programs and their corresponding correct version for the training. Given a set of pairs (`buggy program` p, `fixed version` p'), we first start by filtering out syntactically incorrect code. Next, we instrument the code of both p and p' to get the traces on a line level as defined in Definition 2.2. Then, we execute the instrumented programs while automatically excluding programs terminated with a runtime exception as it is out of the scope of our approach since we are interested in divergence points caused by the difference in the value of the program's state as defined in Definition 2.1. Finally, we compare the trace of each successfully executed buggy program to the trace of its correct version to keep only the buggy programs having a trace that has a divergence point from the trace of the correct version. In some programs, the traces are non-deterministic,
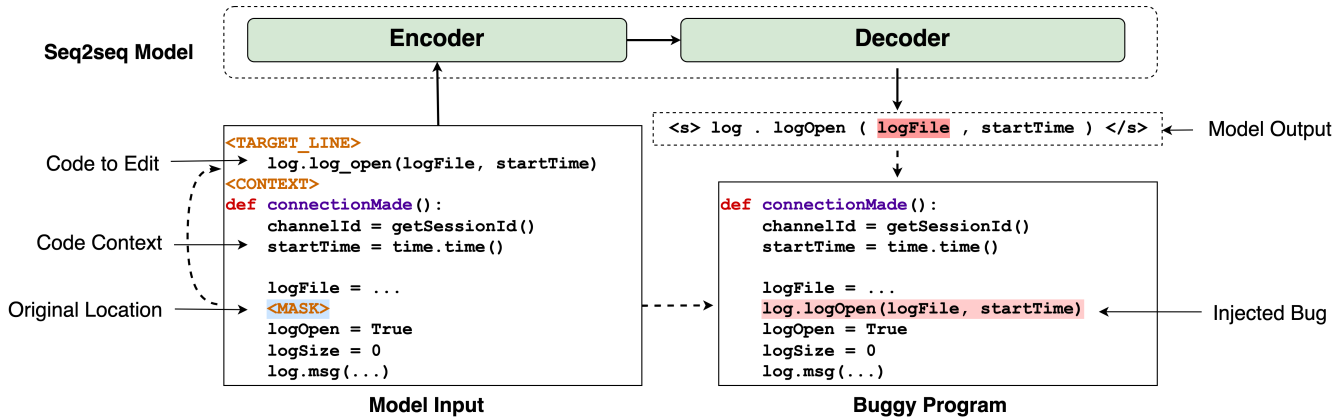
Figure 3: Workflow of bug injection.

e.g., due to the existence of a variable or expression that depends on a random generator or a variant factor (like time). To determine those non-deterministic traces, we execute the original correct code (of the program p') using the same test case at two different timestamps. If the two executions produce different traces, we consider it non-deterministic and drop it. We drop non-deterministic instances because the trace between the original and buggy codes might differ because of that random statement, even if it is not the bug's location. After getting the correct and buggy code trace, we determine the point of divergence by doing a diff between the two traces (correct and buggy). The first location of difference is considered to be the point of divergence. The program's trace ends at the determined point of divergence, and any collected trace after that point is removed. At the end of this step, we get a list of data points (`buggy trace` t, `desired state` s, `buggy code` c, `correct code` c'). To get the desired state, we first locate the point of divergence and copy the state of the correct program at the point of divergence, which becomes the desired state.

Prediction phase. In this phase, the developer executes their code and gets the traces. Then, upon inspecting the trace, the developer locates the divergence point and gives the corresponding desired state. At the end of this step, we get a prediction data point (`buggy trace` t, `desired state` s, `buggy code` c).

2.6 Constructing and Formatting Input

After getting the traces of the programs, we construct the input data for the neural model. The input formatting module takes a data point and then transforms it into a format compatible with the model's input layer. Similar to most previous work on program repair [31], we assume that the location where the code needs to be fixed is known, as this location can be obtained via existing techniques [62].

Definition 2.5 (Formatted data point). a formatted training data point is a pair (`source, target`) where the `source` is the input given to the model, and the `target` is the ground truth for what the model should predict. The input `source` and `target` are defined as follows:

- Source: concatenation of the following modalities separated by the special tokens (< .. >).
  - <BUGGY_LINE> the buggy target line l_b
  - <INITIAL_STATE> initial state s_{init} of the program p
  - <LINE> executed line of code <STATE> the programs state
  - <DESIRED_STATE> the desired or expected correct state
  - <CONTEXT> source code of the program
- Target:
  - <START> the fix of the given buggy line <END>

In prediction phase, a data point consistent of `source` only. We construct the formatted data point from elements in the data point resulting from step two.

Figure 4 shows an example of the input format that follows the format of definition 2.5, including the different modalities. The special separators <..> inserted between the modalities are in bold text. The red dashed line surrounds the buggy line, while the green dashed line delimits the program's state, followed by the desired state below (blue dashed lines). We also provide the context as the last modality. In the example of Figure 4, we represent and give the model the last three states of the buggy code trace. The trace size in terms of the number of states is a configurable parameter in our approach, and it represents the number of traced lines before the divergence point (included). Furthermore, we only include variables that get changed from one state to another. For example, we did not mention the variable `start_port` in the last state because there was no change in its value from the second state to the last one. All the input modalities together form the `source` input. The other part is the `target` delimited by the start and end tokens. The target is the ground truth used to evaluate the model's prediction.

2.7 Deep Learning Model

Once we format the input, we tokenize it using BPE tokenizer (explained in section 2.3). The BPE tokenizer has an encoding function that encodes a string into a list of tokens Ids. A token Id is a number that uniquely encodes and designates a token in the vocabulary of the BPE tokenizer. Figure 4 shows the first tokens Ids resulting from encoding the source input through BPE tokenizer. The BPE
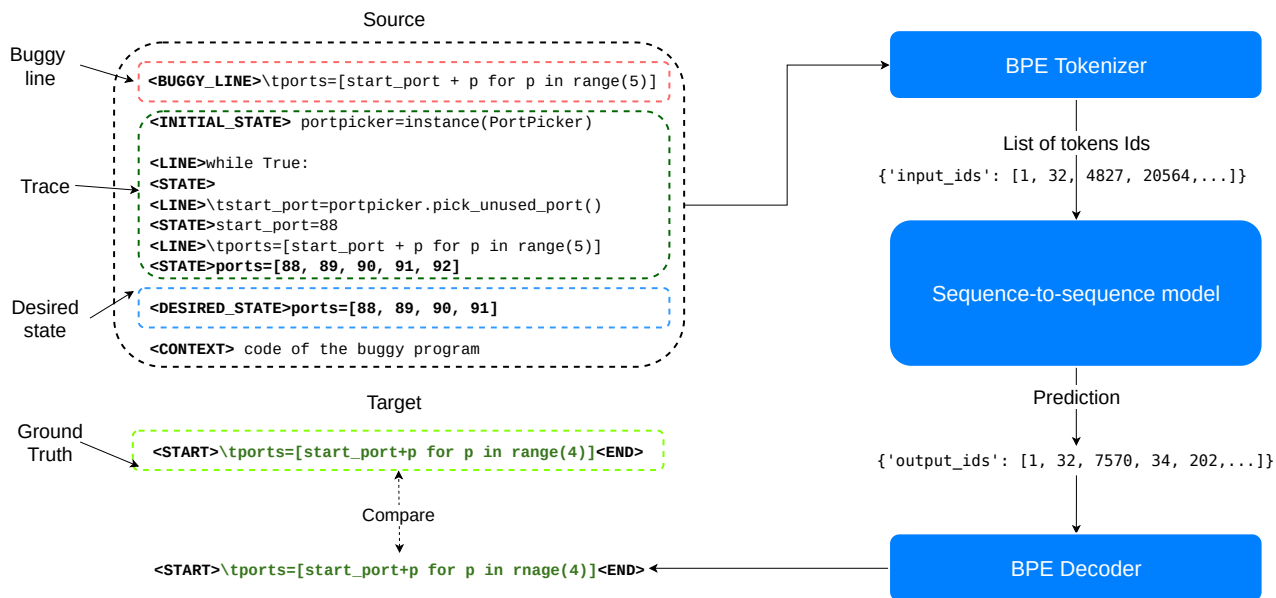
Figure 4: Example of input/output format of the neural model (<..> denote special tokens).

tokenizer also has a decoding function that decodes a list of tokens Ids into a string, as shown in the example of Figure 4.

After tokenization and encoding, the final step in our approach is training the neural model. Our model is a sequence-to-sequence Transformer based on CodeT5, explained in section 2.3. At each training step, the model receives a batch of data points, predicts the corresponding fixes and gets evaluated against the ground truth targets by calculating the loss and then updating the weights accordingly. Given a prediction data point, the model outputs a sequence of tokens Ids, which the BPE then decodes to construct a string corresponding to the predicted fix.

3 EVALUATION

To evaluate the approach, we address the following research questions:

- RQ1: How effective is TraceFixer at fixing bugs?
- RQ2: How does TraceFixer compare to a baseline approach that takes only the code, but no traces, as its input?
- RQ3: How effective is a multi-task model trained to handle both trace-guided repair and code-only repair?
- RQ4: How does the effectiveness of TraceFixer vary across different types of bugs?

3.1 Experimental Setup

3.1.1 Datasets. We use three existing datasets that offer real-world code and bugs in Python.

PySStuBs. This dataset [27] is an extensive collection of single statement bugs in popular open-source Python projects. We use this dataset in two ways. First, we train the bug injector described in Section 2.4 on the PySStuBs bugs. Second, we sample 20 bugs from PySStuBs to assess whether TraceFixer can fix real-world bugs.

The reason for focusing on 20 bugs is that we have to manually set up executable tests that trigger each bug. The setup is non-trivial as each project in PySStuBs has its own requirements and many are no longer compatible with the latest versions of packages and even Python. We made sure these 20 bugs do not overlap with the training dataset.

CodeNet. This dataset [48] is a set of algorithmic challenges where participants try to solve a particular problem and submit their code to the platform to get automatically evaluated against test cases. We use 3,200 problems, where we take three submissions from each problem, adding up to 9,600 different programs. Since the CodeNet dataset does not provide buggy versions of the programs, we automatically inject bugs (Section 2.4). For each program, the bug injector injects three different bugs at every line of code, which yields 850K buggy programs. Following the steps described in Section 2.5 we end up with 103K data points. We keep 500 of them for testing and use the rest to train the neural model.

TheAlgorithms. This dataset[2] is a collection of programs that solve algorithmic problems, along with test cases to validate their correctness. The problems cover various domains, such as arithmetic and algebra, neural networks, and web programming. Similar to the CodeNet dataset, we inject bugs into the given programs with a rate of five variants per line, which yields 120K buggy programs. From the resulting 8.5K datapoints, we keep 300 for testing and use the rest for training.

3.1.2 Metrics. We measure the effectiveness of TraceFixer by checking if the suggested fixes match the known ground-truth fixes. A suggested fix counts as correct if and only if it exactly matches the ground truth, except for unnecessary whitespace. Similar to

---

[2]https://github.com/TheAlgorithms/Python

prior work on learning-based repair [10, 36], we query the model to produce multiple fix suggestions using beam search. We report how often the model finds the correct fix among the unique top-k suggestions, or short UTOPk. The UTOPk metric is a variant of the usual top-k accuracy that ignores duplicate predictions and keeps querying the model until we find k unique fix suggestions. To determine duplicates, we first start by removing unnecessary tokens from the predicted line of code. For example, a = a + 1 and a=a+1 are duplicates because the only difference is unnecessary whitespace.

3.1.3 Baseline. The key contribution of this work is to introduce execution traces into learning-based repair. To assess the benefits of this idea, we compare TraceFixer to a baseline model that takes only the buggy source code as its input, called *Code-only*. Conceptually, this model follows the source code-only, sequence-to-sequence approach proposed in recent learning-based repair techniques [8, 10, 36]. We implement the baseline as a stripped-down variant of TraceFixer, which differs from the full TraceFixer approach only by not receiving any execution traces as the input. This code-only setting is similar to prior work, MODIT [8]—both use similar input formatting. A possible alternative would be to compare to the implementation of prior code-only repair techniques directly. We decided against that setup (i) because prior tools focus on Java instead of Python and (ii) because it would bias the results due to details of the techniques that are not directly related to the question of whether traces are helpful for learning-based repair.

3.1.4 Implementation and Hardware. We implement the different components of TraceFixer into a Python-based tool. To instrument and collect traces, we use Python standard modules such as SetTrace and Inspect alongside other community packages like beeprint.[3] The neural models are based on CodeT5-small shared through hugging face api.[4] We conduct our experiments on two GPU machines. The first machine is equipped with an NVIDIA Tesla P100 (16GB), while the second has an NVIDIA Tesla V100 (32GB). To save time, we split our experiments across the two machines, which allows us to train one epoch in about 2.5 hours on the first machine and in about 1.5 hours on the second machine.

3.2 RQ1: Bug Fixing Effectiveness

To investigate the effectiveness of TraceFixer in program repair, we evaluate our approach on three different datasets: CodeNet, TheAlgorithms and 20 real bugs sampled from PySStuBs. For the CodeNet and TheAlgorithms datasets, we finetune a separate CodeT5-small model for each dataset and then evaluate on a held-out test set from the same dataset. To evaluate TraceFixer on the 20 real bugs, we use the model trained on CodeNet.

The results of the evaluation on the test sets are shown in Table 1. The table shows that TraceFixer achieves high accuracy across the three setups. For example, considering only the top-most prediction by the model fixes 63%, 57%, and 10% of the bugs, while the predictions in the UTOP10 contain correct fixes for 87%, 82%, and 50% of all bugs. The fixing accuracy is the lowest on the set of real bugs, which we attribute to three reasons. First, the model is trained on a

---

[3]https://github.com/panyanyany/beeprint
[4]https://huggingface.co/Salesforce/codet5-small

different dataset (CodeNet) and then applied to the real-world bugs, which may cause a shift of distribution [21]. Second, the real bugs are extracted from open-source projects that are more complicated than the solutions to algorithmic problems in the CodeNet and TheAlgorithms datasets. Finally, the real-world bugs occur in code that calls other modules and functions within the same project or other packages. In contrast, the CodeNet code mainly uses standard Python functions and data structures.

For a more detailed look into the 20 real bugs, Table 2 shows all of them alongside the fixes and the effectiveness of the model. We can see that the TraceFixer is able to fix different kinds of problems, including control flow-related bugs (e.g., example 4), data structure-related bugs (e.g., examples 8 and 10), and various kinds of incorrect expressions (e.g., examples 1 and 14). The cases in which TraceFixer cannot find a fix often are due to the out-of-vocabulary problem [28], such as in examples 2, 3 and 18. Out-of-vocabulary here means that the fix requires some tokens that do appear neither in the provided code nor the execution trace, and unless the model has seen such a fix during training, it typically cannot predict the desired tokens.

> **Finding 1:** Throughout the different evaluation setups, TraceFixer effectively predicts the correct fix among the UTOP10 in 50%–87% of the cases.

3.3 RQ2: Comparison with Code-Only Approach

The key contribution of TraceFixer is to use an execution trace and the desired state as an additional input to a repair model. To assess the impact of this contribution over prior work, we compare our approach to the code-only baseline. Table 1 shows that across the three evaluation setups, TraceFixer improves over the code-only baselines in terms of UTOP K accuracy. For example, in the CodeNet evaluation setup, TraceFixer improves with 13% over the UTOP 1. Using three different datasets suggests that the improvement over the code-only approach is independent of the dataset.

Being powered by the trace, our approach provides the model with more information to fix the bug, which aligns with the analogy of the developer's way of debugging. While the code-only model still performs well on the three evaluation datasets, it has limited knowledge about the desired change and thus does not predict the most likely fix in many cases. Our approach mitigates the limitation of code-only models by adding traces and the desired states as data modalities, thus giving more guidance to perform the repair task.

Figure 5 shows the Venn diagram of the overlap in the correctly fixed instances between TraceFixer and code-only across the three datasets. For example, the first part of Figure 5 shows that TraceFixer predicts correct fixes among the UTOP10 for almost all the instances that code-only can fix with UTOP10 predictions. Similarly, only 2% of the TheAlgorithms test set is fixed correctly by the code-only model and not fixed by TraceFixer. Finally, in the real bugs, the set of fixed instances by the code-only model is strictly included in the set of instances fixed by TraceFixer. In summary, our approach fixes more than 97% of the instances that the code-only baseline can fix across the three datasets. Of course, TraceFixer fixes other instances that the code-only model cannot fix.

Table 1: UTOPk accuracy on different datasets.

Dataset	Models	UTOP1(%)	UTOP3(%)	UTOP5(%)	UTOP10(%)
CodeNet	TraceFixer	**63**	**78**	**82**	**87**
	Code-only	50	65	69	78
TheAlgorithms	TraceFixer	**57**	**75**	**77**	**82**
	Code-only	49	64	66	69
20 Real bugs	TraceFixer	10	**25**	**30**	**50**
	Code-only	10	15	25	30
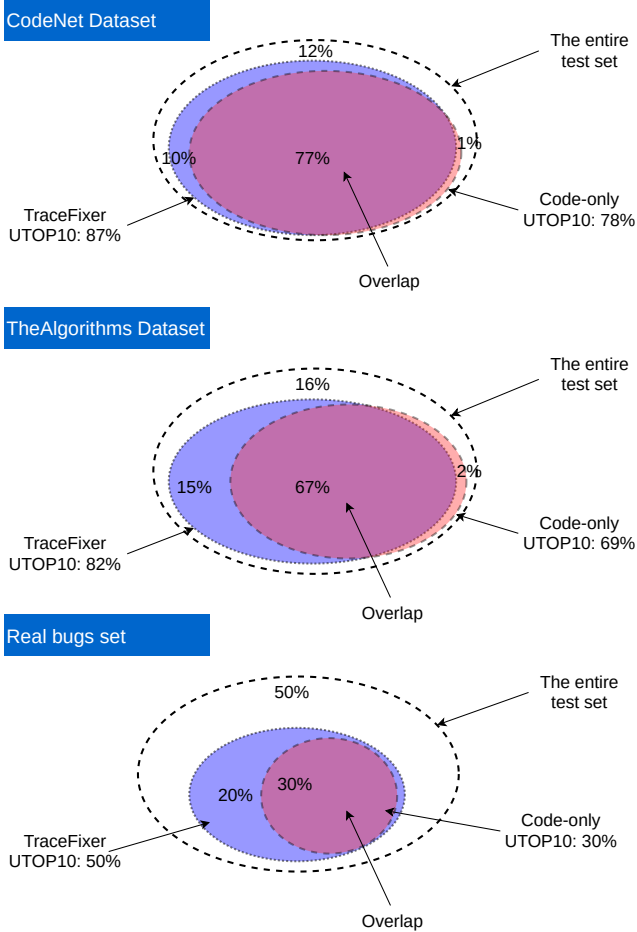
Figure 5: Overlap between TraceFixer and code-only fixed instances across the three datasets (Venn diagram).

3.4 RQ3: Combining Trace-Guided and Code-Only Learning

Multitask learning is a deep-learning training setup in which one model learns to perform multiple tasks by training on data encoding multiple tasks. The goal and intuition behind multitask learning is that the model gains a better understanding by looking at the data from different perspectives that reflect the specificity of the task. More formally, the deep learning model trains to find an internal representation or embedding of data that generalizes to both tasks. There has been much work on multitask learning recently, and it has many applications in computer vision, NLP and software analysis.

In this experiment, we want to explore combining the baseline approach with our approach in one model through the multitask learning setup. Therefore, we use the CodeNet dataset for training. The dataset comprises inputs without trace modality (code-only) and inputs with trace modality (our approach). The two types of inputs corresponding to the two tasks are distinguished using two prefixes: <FFT> (fix from trace) and <FFC> (fix from code). When querying the multitask model, we first pass a trace-based data point. Then, if the model cannot fix the bug from trace-based input, we pass the equivalent code-only data point. In other words, we ask the model to fix from the trace. Then, if the predicted fixes are incorrect, we ask the model again to fix from code only. Since the multitask model performs twice the UTOPk for each query, we use UTOPk with k in (2, 6, 10, 20) to ensure fairness with other unitask models.

In Table 3, we compare the UTOPk accuracy of the multitask model with the single-task models. The table shows a slight improvement of the multitask over our default TraceFixer. From one side, we can understand that training in a multitask setup improves overall performance. Conversely, we can see that code-only and trace-based program repair can be complementary. The approach can benefit from reasoning about the code only when the trace is unavailable or does not help fix the code.

Finding 2: Throughout the different evaluation setups, TraceFixer finds the correct fix more often than the code-only baseline. Furthermore, almost all the instances (more than 97%) fixed by code-only model are also fixed by TraceFixer

Finding 3: Training TraceFixer in a multitask setup, where the model can use code-only mode or trace-based mode, improves over the single-task variant of TraceFixer and encodes both TraceFixer and code-only baseline in one model.

Table 2: The list of 20 real bugs used in RQ2.

Id	Buggy line	Ground truth and Model's prediction	Project	Correct?
1	`num_lines = size_log_area / 2`	GT/Model: `num_lines = size_log_area // 2`	MycroftAI	Y
2	`return ((x + y) == 1) * 1.0`	GT: `return 1 - abs(x + y - 1)` Model: `return ((x + y) == 0) * 1.0`	keras	N
3	`return rs.rand(3)`	GT: `return rs.rand(3).tolist()` Model: `return int(rs.rand(3))`	pandas	N
4	`if not columns or len(columns) == 9:`	GT/Model: `if not columns or len(columns) == 0:`	great-expectations	Y

5	`schedule_json["interval"] = query.`
	`    old_schedule`
		GT: `schedule_json["interval"] = int(query.`
		`    old_schedule)`
		Model: `schedule_json["interval"] = query.`
		`    new_schedule`
			redash	N

6	`step_size = step_size_scaling * n**(1/4.)`	GT: `step_size = step_size_scaling / n**(1/4.)` Model: `step_size = step_size_scaling * n**(1/2.)`	pymc	N
7	`limit = parse_integer(request, "limit", 100)`	GT/Model: `limit = parse_integer(request, "limit",` `    0)`	matrix-org/synapse	Y
8	`ports = [start_port + p for p in range(5)]`	GT/Model: `ports = [start_port + p for p in range(` `    4)]`	deepmind/pysc2	Y
9	`self._ports = _pick_unused_ports(1 + self.` `    _num_players * 2)`	GT/Model: `self._ports = _pick_unused_ports(self.` `    _num_players * 2)`	deepmind/pysc2	Y
10	`_transfer_block(src, dst.res4, ['4a'] + ['4b%` `    d' % i for i in range(1, 23)])`	GT/Model: `_transfer_block(src, dst.res4, ['4a'] +` `    ['4b%d' % i for i in range(1, 24)])`	chainer	Y
11	`return 0.5*(ts-bs)`	GT: `return 0.5 - bs` Model: `return 0.5*(ts+bs)`	3b1b/manim	N
12	`self.stdinlogOpen = True`	GT/Model: `self.stdinlogOpen = False`	cowrie	Y
13	`self.terminal.stdinlog_open=True`	GT/Model: `self.terminal.stdinlogOpen=True`	cowrie	Y

14	`progress_total = self.parameter_depth * len(`
	`    self.http_methods) * (`
	`    2 + len(self.integration_response_codes) +`
	`    len(self.method_response_codes)) - 1`
		GT/Model: `progress_total = self.parameter_depth *`
		`    len(self.http_methods) * (`
		`    2 + len(self.integration_response_codes) + len(`
		`    self.method_response_codes))`
			Miserlou/Zappa	Y

15	`outputs = outputs * num_units**0.5`	GT: `outputs = outputs / num_units**0.5` Model: `outputs = outputs * num_units**0.5 / 2`	Kyubyong/ transformer	N
16	`drive_start = drive_letter(path)`	GT/Model: `drive_start = drive_letter(start)`	numba	Y
17	`value = utils.force_type(gy[0].dtype, self.` `    value)`	GT: `value = utils.force_type(x.dtype, self.value)` Model: `value = utils.force_type(gy[1].dtype,` `    value)`	chainer	N
18	`output = tout.readlines()`	GT: `output = tout.read().splitlines()` Model: `output = list(tout.readlines())`	salt	N
19	`join.shared_port = self._ports.pop()`	GT: `join.shared_port = 0` Model: `join.shared_port = self._ports.pop(0)`	deepmind/pysc2	N
20	`ret['changes']['interface'] = ''.join(diff)`	GT: `ret['changes']['interface'] = '\n'.join(diff)` Model: `ret['changes']['interface'] = (' ').join(` `    diff)`	salt	N

Table 3: Accuracy of a multitask variant of TraceFixer compared with single-task variants (CodeNet dataset).

Models	UTOP2(%)	UTOP6(%)	UTOP10(%)	UTOP20(%)
Multitask	71	**86**	**88**	**90**
TraceFixer	**72**	82	87	88
Code-only	64	73	78	85

Table 4: Accuracy of the approach on different bug types.

Bug type	UTOP1	UTOP3	UTOP5	UTOP10
Arithmetic	75	**85**	**90**	**90**
Varmisuse	75	**85**	**90**	**90**
Functions	35	55	60	65
Data structures	**80**	80	80	80
Control flow	60	80	80	80

3.5 RQ4: Effectiveness Depending on the Kind of Bug

In this experiment, we want to explore further how well the approach performs on different kinds of bugs. For that, we first categorize the test data sampled from CodeNet data into five categories. Arithmetic bugs are instances where a numeral operand or operator is wrong in an arithmetic expression, e.g., `speed = distance * time` instead of `speed = distance / time`. The Varmisuse category represents bugs where the wrong variable is used in a statement. For example, row 16 of Table 2 represents an example of a varmisuse bug. The category Functions include missing function call on an expression, calling the wrong function, or calling a function incorrectly. For example, in the following code: `graph = sorted(graph.items())`, the function `sorted` is called in the wrong way because it is missing the initialization of the argument key, and thus the code should be: `graph = sorted(graph.items(), key=lambda x:len(x[1]))`. The bugs related to data structures comprise bugs where a data structure is wrongly manipulated, for example, `list.push(0)` instead of `list.pop(0)`. Finally, control flow bugs occur in control flow statements such as `if`, `for` and `while`. Each category has 20 instances. The second step in our experiment is to use the model trained on the CodeNet dataset to evaluate each category. Similar to previous experiments, we calculate UTOP k for each category where k is 1, 3, 5, or 10.

Table 4 summarizes the model's performance across the five categories. The model performance on different categories is close to the overall performance of TraceFixer except for the Functions category. The drop in performance on the Functions category is mainly because of the trace's lack of guidance to perform the proper fix. In the previous example, adding the key argument will change the sorting key of the elements of a list or iterable. The sorting is based on the values of the second element of each tuple in a list of tuples. In the end, the difference between the divergence state of the code and the desired state is the order of elements in the list which is difficult to translate into the addition of the specific new tokens: `key=lambda x:len(x[1])`.

Finding 4: Our approach is effective across different kinds of bugs. It is still challenging for TraceFixer to fix bugs where the edit requires out-of-vocabulary tokens or needs documentation on how to use a specific function or API.

3.6 Threats to Validity

Generalization to other languages. Our approach by design is independent of the target programming language. However, our implementation is entirely in Python and for Python programs. We targeted Python because there needs to be more work on Program repair that evaluates on Python code. While our approach is unrelated to the specifications of any language, it is still interesting to evaluate on other languages to investigate whether the same results can be replicated.

Bug localization and tracing. Our approach works under the hypothesis that the developer can localize the buggy line and give the desired state at the point of divergence through debugging or other tools. While there is work on localizing bugs, providing the desired state is still an issue. It might take much work for the developer to accurately provide the desired state in complicated programs or programs manipulating large data structures.

4 RELATED WORK

Traditional approaches to program repair. Program repair and fixing has been extensively studied in the literature [17, 18, 38]. Traditional program repair often follows a pipeline of steps. It starts with fault localization [29, 34, 50, 62, 64, 67], followed by searching fix candidates based on various search strategies [25, 59, 61]. The fixes are then checked against correctness criteria (e.g., executing tests). Generating fix candidates is often the most expensive part as it requires searching in a combinatorial space of the program fixes. While there exist strategies to reduce the search space by limiting the fix templates, i.e., using only insertions, deletions, or existing code snippets appeared in the program [20], but it may not be expressive enough to represent the correct fixes. TraceFixer automates the search process *over the entire vocabulary of the program space* by producing fixes directly via an efficient inference pass (using GPUs).

Learning-based program repair. Machine learning has been widely used to localize bugs [3, 29, 33, 35, 49] and generate program fixes [2, 6, 13, 15, 24, 49, 53, 66, 69], the two critical steps in automated program repair. In particular, to localize bugs, Lou et al. [35] used tests *coverage* and employed a gated graph neural network on ASTs of buggy code and test cases. DeepBugs [47] embeds single expression to predict 3 types of bugs: swapped function arguments, incorrect binary operator, and incorrect operand in a binary operation. Li et al. [33] learns a convolution network based on the code coverage, statement dependency, and stack traces to predict bug locations. Qi et al. [49] exploits bug reports to localize bugs. To generate repairs, Devlin et al. [12] and Vasic et al. [54] leverage a pointer network to predict the fixes for a restricted set of bug classes. DeepFix [19] learns an attention-based sequence-to-sequence network to localize and generate fixes for C programs. Codit [6] learns a tree-based network to better capture the structural changes for

program fix. CoditT5 [66] pretrains a large language model to explicitly learn the general edit operations, and finetune it on three downstream code editing tasks. Unfortunately, none of the existing program repair approaches consider dynamic program traces, which provide an important hint to repair programs. As a result, the class of bugs that TraceFixer can repair is much broader than prior works.

Learning representations for programming tasks. With abundant open-source software and computing power available, deep neural networks have been broadly used to automate software development processes [45]. Examples have these applications include type inference [22, 43, 46, 63], program generation [5, 9, 11, 32], code summarization [52, 56], bug/vulnerability detection [7, 14, 16, 42], and code clone detection [14, 37, 44]. It has been shown that incorporating traces benefits learning program representations for downstream applications that relies on understanding program semantics [26, 40, 42, 43, 57]. We demonstrate that feeding the trace divergence to neural networks can significantly improve its effectiveness of repairing program bugs (3.2).

5 CONCLUSIONS

In this paper, we present TraceFixer, a new approach for learning-based program repair that incorporates the execution trace into the input of the deep learning model. The evaluation on multiple datasets shows that TraceFixer is effective at fixing bugs, improves over code-only baseline approaches, and allows to repair more real bugs.

REFERENCES

[1] Wasi Uddin Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified Pre-training for Program Understanding and Generation. *CoRR* abs/2103.06333 (2021). arXiv:2103.06333 https://arxiv.org/abs/2103.06333

[2] Miltiadis Allamanis, Earl T Barr, Christian Bird, and Charles Sutton. 2014. Learning natural coding conventions. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. 281–293.

[3] Miltiadis Allamanis, Marc Brockschmidt, and Mahmoud Khademi. 2017. Learning to represent programs with graphs. *arXiv preprint arXiv:1711.00740* (2017).

[4] Miltiadis Allamanis, Henry Jackson-Flux, and Marc Brockschmidt. 2021. Self-Supervised Bug Detection and Repair. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 27865–27876.

[5] Saikat Chakraborty, Toufique Ahmed, Yangruibo Ding, Premkumar Devanbu, and Baishakhi Ray. 2022. NatGen: Generative pre-training by" Naturalizing" source code. In *Proceedings of the 2022 30th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*.

[6] Saikat Chakraborty, Yangruibo Ding, Miltiadis Allamanis, and Baishakhi Ray. 2020. Codit: Code editing with tree-based neural models. *IEEE Transactions on Software Engineering* (2020).

[7] Saikat Chakraborty, Rahul Krishna, Yangruibo Ding, and Baishakhi Ray. 2021. Deep learning based vulnerability detection: Are we there yet. *IEEE Transactions on Software Engineering* (2021).

[8] S. Chakraborty and B. Ray. 2021. On Multi-Modal Learning of Editing Source Code. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE Computer Society, Los Alamitos, CA, USA, 443–455. https://doi.org/10.1109/ASE51524.2021.9678559

[9] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).

[10] Zimin Chen, Steve Kommrusch, Michele Tufano, Louis-Noël Pouchet, Denys Poshyvanyk, and Martin Monperrus. 2019. SequenceR: Sequence-to-Sequence Learning for End-to-End Program Repair. *IEEE TSE* (2019).

[11] Matteo Ciniselli, Nathan Cooper, Luca Pascarella, Denys Poshyvanyk, Massimiliano Di Penta, and Gabriele Bavota. 2021. An empirical study on the usage of BERT models for code completion. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 108–119.

[12] Jacob Devlin, Jonathan Uesato, Rishabh Singh, and Pushmeet Kohli. 2017. Semantic code repair using neuro-symbolic transformation networks. *arXiv preprint arXiv:1710.11054* (2017).

[13] Elizabeth Dinella, Todd Mytkowicz, Alexey Svyatkovskiy, Christian Bird, Mayur Naik, and Shuvendu K Lahiri. 2021. Deepmerge: Learning to merge programs. *arXiv preprint arXiv:2105.07569* (2021).

[14] Yangruibo Ding, Luca Buratti, Saurabh Pujar, Alessandro Morari, Baishakhi Ray, and Saikat Chakraborty. 2022. Towards Learning (Dis)-Similarity of Source Code from Program Contrasts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 6300–6312.

[15] Yangruibo Ding, Baishakhi Ray, Premkumar Devanbu, and Vincent J Hellendoorn. 2020. Patching as translation: the data and the metaphor. In *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 275–286.

[16] Yangruibo Ding, Sahil Suneja, Yunhui Zheng, Jim Laredo, Alessandro Morari, Gail Kaiser, and Baishakhi Ray. 2022. VELVET: a noVel Ensemble Learning approach to automatically locate VulnErable sTatements. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 959–970.

[17] Luca Gazzola, Daniela Micucci, and Leonardo Mariani. 2018. Automatic software repair: A survey. In *Proceedings of the 40th International Conference on Software Engineering*. 1219–1219.

[18] Claire Le Goues, Michael Pradel, and Abhik Roychoudhury. 2019. Automated program repair. *Commun. ACM* 62, 12 (2019), 56–65.

[19] Rahul Gupta, Soham Pal, Aditya Kanade, and Shirish Shevade. 2017. Deepfix: Fixing common c language errors by deep learning. In *Thirty-First AAAI conference on artificial intelligence*.

[20] Mark Harman. 2010. Automated patching techniques: the fix is in: technical perspective. *Commun. ACM* 53, 5 (2010), 108–108.

[21] Jingxuan He, Luca Beurer-Kellner, and Martin T. Vechev. 2022. On Distribution Shift in Learning-based Bug Detectors. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 8559–8580. https://proceedings.mlr.press/v162/he22a.html

[22] Vincent J Hellendoorn, Christian Bird, Earl T Barr, and Miltiadis Allamanis. 2018. Deep learning type inference. In *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.

[23] Vincent J. Hellendoorn and Premkumar Devanbu. 2017. Are Deep Neural Networks the Best Choice for Modeling Source Code? *(ESEC/FSE 2017)*. Association for Computing Machinery, New York, NY, USA, 763–773. https://doi.org/10.1145/3106237.3106290

[24] Faria Huq, Masum Hasan, Md Mahim Anjum Haque, Sazan Mahbub, Anindya Iqbal, and Toufique Ahmed. 2022. Review4Repair: Code review aided automatic program repairing. *Information and Software Technology* 143 (2022), 106765.

[25] Jiajun Jiang, Yingfei Xiong, Hongyu Zhang, Qing Gao, and Xiangqun Chen. 2018. Shaping program repair space with existing patches and similar code. In *Proceedings of the 27th ACM SIGSOFT international symposium on software testing and analysis*. 298–309.

[26] Xin Jin, Kexin Pei, Jun Yeon Won, and Zhiqiang Lin. 2022. SymLM: Predicting Function Names in Stripped Binaries via Context-Sensitive Execution-Aware Code Embeddings. In *2022 ACM SIGSAC Conference on Computer and Communications Security*.

[27] Arthur V. Kamienski, Luisa Palechor, Cor-Paul Bezemer, and Abram Hindle. 2021. PySStuBs: Characterizing Single-Statement Bugs in Popular Open-Source Python Projects. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. 520–524. https://doi.org/10.1109/MSR52588.2021.00066

[28] Rafael-Michael Karampatsis, Hlib Babii, Romain Robbes, Charles Sutton, and Andrea Janes. 2020. Big code != big vocabulary: open-vocabulary models for source code. In *ICSE '20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*, Gregg Rothermel and Doo-Hwan Bae (Eds.). ACM, 1073–1085. https://doi.org/10.1145/3377811.3380342

[29] Yunho Kim, Seokhyeon Mun, Shin Yoo, and Moonzoo Kim. 2019. Precise learn-to-rank fault localization using dynamic and static features of target programs. *ACM Transactions on Software Engineering and Methodology* (2019).

[30] Claire Le Goues, ThanhVu Nguyen, Stephanie Forrest, and Westley Weimer. 2012. GenProg: A Generic Method for Automatic Software Repair. *IEEE Trans. Software Eng.* 38, 1 (2012), 54–72.

[31] Claire Le Goues, Michael Pradel, and Abhik Roychoudhury. 2019. Automated program repair. *Commun. ACM* 62, 12 (2019), 56–65. https://doi.org/10.1145/3318162

[32] Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. *arXiv preprint arXiv:2203.07814* (2022).

[33] Yi Li, Shaohua Wang, and Tien N Nguyen. 2021. Fault localization with code coverage representation learning. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 661–673.

[34] Yiling Lou, Ali Ghanbari, Xia Li, Lingming Zhang, Dan Hao, and Lu Zhang. 2019. Can Automated Program Repair Refine Fault Localization? *arXiv preprint*

arXiv:1910.01270 (2019).

[35] Yiling Lou, Qihao Zhu, Jinhao Dong, Xia Li, Zeyu Sun, Dan Hao, Lu Zhang, and Lingming Zhang. 2021. Boosting coverage-based fault localization via graph-based representation learning. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 664–676.

[36] Thibaud Lutellier, Hung Viet Pham, Lawrence Pang, Yitong Li, Moshi Wei, and Lin Tan. 2020. CoCoNuT: combining context-aware neural translation models using ensemble for program repair. In *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, USA, July 18-22, 2020*, Sarfraz Khurshid and Corina S. Pasareanu (Eds.). ACM, 101–114. https://doi.org/10.1145/3395363.3397369

[37] Nikita Mehrotra, Navdha Agarwal, Piyush Gupta, Saket Anand, David Lo, and Rahul Purandare. 2021. Modeling functional similarity in source code with graph-based Siamese networks. *IEEE Transactions on Software Engineering* (2021).

[38] Martin Monperrus. 2018. Automatic software repair: a bibliography. *ACM Computing Surveys (CSUR)* 51, 1 (2018), 1–24.

[39] Yu Nong, Yuzhe Ou, Michael Pradel, Feng Chen, and Haipeng Cai. 2022. Generating Realistic Vulnerabilities via Neural Code Editing: An Empirical Study. In *ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*.

[40] Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. 2021. Show Your Work: Scratchpads for Intermediate Computation with Language Models. *arXiv preprint arXiv:2112.00114* (2021).

[41] Jibesh Patra and Michael Pradel. 2021. Semantic bug seeding: a learning-based approach for creating realistic bugs. In *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23-28, 2021*, Diomidis Spinellis, Georgios Gousios, Marsha Chechik, and Massimiliano Di Penta (Eds.). ACM, 906–918. https://doi.org/10.1145/3468264.3468623

[42] Jibesh Patra and Michael Pradel. 2022. Nalin: Learning from Runtime Behavior to Find Name-Value Inconsistencies in Jupyter Notebooks. In *2022 IEEE 31st International Conference on Software Engineering (ICSE)*.

[43] Kexin Pei, Jonas Guan, Matthew Broughton, Zhongtian Chen, Songchen Yao, David Williams-King, Vikas Ummadisetty, Junfeng Yang, Baishakhi Ray, and Suman Jana. 2021. StateFormer: fine-grained type recovery from binaries using generative state modeling. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 690–702.

[44] Kexin Pei, Zhou Xuan, Junfeng Yang, Suman Jana, and Baishakhi Ray. 2021. TREX: Learning Execution Semantics from Micro-Traces for Binary Similarity. In *2021 IEEE Symposium on Security and Privacy*.

[45] Michael Pradel and Satish Chandra. 2022. Neural software analysis. *Commun. ACM* 65, 1 (2022), 86–96. https://doi.org/10.1145/3460348

[46] Michael Pradel, Georgios Gousios, Jason Liu, and Satish Chandra. 2020. Type-writer: Neural type prediction with search-based validation. In *28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.

[47] Michael Pradel and Koushik Sen. 2018. Deepbugs: A learning approach to name-based bug detection. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 1–25.

[48] Ruchir Puri, David S. Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir R. Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, Shyam Ramji, Ulrich Finkler, Susan Malaika, and Frederick Reiss. 2021. CodeNet: A Large-Scale AI for Code Dataset for Learning a Diversity of Coding Tasks. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*, Joaquin Vanschoren and Sai-Kit Yeung (Eds.). https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/a5bfc9e07964f8dddeb95fc584cd965d-Abstract-round2.html

[49] Binhang Qi, Hailong Sun, Wei Yuan, Hongyu Zhang, and Xiangxin Meng. 2021. DreamLoc: A Deep Relevance Matching-Based Framework for bug Localization. *IEEE Transactions on Reliability* (2021).

[50] Ripon K Saha, Matthew Lease, Sarfraz Khurshid, and Dewayne E Perry. 2013. Improving bug localization using structured information retrieval. In *2013 28th IEEE/ACM International Conference on Automated Software Engineering*.

[51] Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. http://aclweb.org/anthology/P/P16/P16-1162.pdf

[52] Ensheng Shia, Yanlin Wangb, Lun Dub, Junjie Chenc, Shi Hanb, Hongyu Zhangd, Dongmei Zhangb, and Hongbin Suna. 2022. On the Evaluation of Neural Code Summarization. In *2022 International Conference on Software Engineering*.

[53] Alexey Svyatkovskiy, Todd Mytkowicz, Negar Ghorbani, Sarah Fakhoury, Elizabeth Dinella, Christian Bird, Neel Sundaresan, and Shuvendu Lahiri. 2021. Merge-BERT: Program Merge Conflict Resolution via Neural Transformers. *arXiv*

[54] Marko Vasic, Aditya Kanade, Petros Maniatis, David Bieber, and Rishabh Singh. 2019. Neural program repair by jointly learning to localize and repair. *arXiv preprint arXiv:1904.01720* (2019).

[55] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*. 6000–6010. http://papers.nips.cc/paper/7181-attention-is-all-you-need

[56] Chong Wang, Xin Peng, Mingwei Liu, Zhenchang Xing, Xuefang Bai, Bing Xie, and Tuo Wang. 2019. A learning-based approach for automatic construction of domain glossary from source code and documentation. In *Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*. 97–108.

[57] Ke Wang and Zhendong Su. 2020. Blended, precise semantic program embeddings. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 121–134.

[58] Yue Wang, Weishi Wang, Shafiq Joty, and Steven C.H. Hoi. 2021. CodeT5: Identifier-aware Unified Pre-trained Encoder-Decoder Models for Code Understanding and Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 8696–8708. https://doi.org/10.18653/v1/2021.emnlp-main.685

[59] Westley Weimer, Zachary P Fry, and Stephanie Forrest. 2013. Leveraging program equivalence for adaptive program repair: Models and first results. In *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 356–366.

[60] Westley Weimer, ThanhVu Nguyen, Claire Le Goues, and Stephanie Forrest. 2009. Automatically Finding Patches Using Genetic Programming. In *International Conference on Software Engineering (ICSE)*. 363–374.

[61] Ming Wen, Junjie Chen, Rongxin Wu, Dan Hao, and Shing-Chi Cheung. 2018. Context-aware patch generation for better automated program repair. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. IEEE, 1–11.

[62] W Eric Wong, Ruizhi Gao, Yihao Li, Rui Abreu, and Franz Wotawa. 2016. A survey on software fault localization. *IEEE Transactions on Software Engineering* 42, 8 (2016), 707–740.

[63] Zhaogui Xu, Xiangyu Zhang, Lin Chen, Kexin Pei, and Baowen Xu. 2016. Python probabilistic type inference with natural language support. In *24th ACM SIGSOFT international symposium on foundations of software engineering*.

[64] Bo Yang, Yuze He, Huai Liu, Yixin Chen, and Zhi Jin. 2020. A Lightweight Fault Localization Approach based on XGBoost. In *2020 IEEE 20th International Conference on Software Quality, Reliability and Security*.

[65] He Ye, Matias Martinez, and Martin Monperrus. 2022. Neural Program Repair with Execution-based Backpropagation. In *ICSE*.

[66] Jiyang Zhang, Sheena Panthaplackel, Pengyu Nie, Junyi Jessy Li, and Milos Gligoric. 2022. CoditT5: Pretraining for Source Code and Natural Language Editing. *arXiv preprint arXiv:2208.05446* (2022).

[67] Mengshi Zhang, Xia Li, Lingming Zhang, and Sarfraz Khurshid. 2017. Boosting spectrum-based fault localization using pagerank. In *26th ACM SIGSOFT International Symposium on Software Testing and Analysis*.

[68] Qihao Zhu, Zeyu Sun, Yuan-an Xiao, Wenjie Zhang, Kang Yuan, Yingfei Xiong, and Lu Zhang. 2021. A syntax-guided edit decoder for neural program repair. In *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23-28, 2021*, Diomidis Spinellis, Georgios Gousios, Marsha Chechik, and Massimiliano Di Penta (Eds.). ACM, 341–353. https://doi.org/10.1145/3468264.3468544

[69] Qihao Zhu, Zeyu Sun, Yuan-an Xiao, Wenjie Zhang, Kang Yuan, Yingfei Xiong, and Lu Zhang. 2021. A syntax-guided edit decoder for neural program repair. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.