

Repairing Bugs in Python Assignments using Large Language Models

Jialu Zhang
Yale University
New Haven, USA
jialu.zhang@yale.edu

José Cambronero
Microsoft
New Haven, USA
jcambronero@microsoft.com

Sumit Gulwani
Microsoft
Redmond, USA
sumitg@microsoft.com

Vu Le
Microsoft
Redmond, USA
levu@microsoft.com

Ruzica Piskac
Yale University
New Haven, USA
ruzica.piskac@yale.edu

Gustavo Soares
Microsoft
Redmond, USA
gsoares@microsoft.com

Gust Verbruggen
Microsoft
Keerbergen, Belgium
gverbruggen@microsoft.com

Abstract—Students often make mistakes on their introductory programming assignments as part of their learning process. Unfortunately, providing custom repairs for these mistakes can require a substantial amount of time and effort from class instructors. Automated program repair (APR) techniques can be used to synthesize such fixes. Prior work has explored the use of symbolic and neural techniques for APR in the education domain. Both types of approaches require either substantial engineering efforts or large amounts of data and training. We propose to use a large language model trained on code, such as Codex, to build an APR system – MMAPR – for introductory Python programming assignments. Our system can fix *both* syntactic and semantic mistakes by combining multi-modal prompts, iterative querying, test-case-based selection of few-shots, and program chunking. We evaluate MMAPR on 286 real student programs and compare to a baseline built by combining a state-of-the-art Python syntax repair engine, BIFI, and state-of-the-art Python semantic repair engine for student assignments, Refactory. We find that MMAPR can fix more programs and produce smaller patches on average.

I. INTRODUCTION

Programming education has grown substantially in popularity in the past decade [1]. A key challenge associated with this growth is the need to provide novice students with efficient and effective learning support. In an ideal world, teaching assistants would monitor students’ learning process, and when students’ code is not correct, they would then help them to derive a correct solution. However, this approach does not scale and educational institutions struggle to find teaching assistants. As a result, there is an interest in developing automated tools that students can use for feedback instead. These tools provide custom repairs for their programming mistakes. The field of automated program repair (APR), which has a long history in the software engineering community [2]–[7], has introduced different approaches [8]–[11] to produce such automated repairs for student mistakes in introductory assignments. Given a buggy student program, the APR system aims to produce a patch that satisfies a specification (typically the instructor-provided test

cases). The patch must also minimize the number of changes made, with the goal of facilitating student learning [11].

Prior automated program repair systems for student programming assignments have generally been implemented using purely symbolic [9]–[12] or purely neural [8], [13] techniques. Symbolic approaches require substantial engineering efforts to develop, typically requiring significant program analysis/repair experience, as well as custom repair strategies tailored to the language domain in which students implement their assignments. Neural approaches mitigate some of the engineering challenges but typically require substantial amounts of data, often leading to specialized use cases for Massive Open Online Courses (MOOCs). Furthermore, these systems are typically tailored to focus exclusively on syntax repair or exclusively on semantic repair. For the latter, the assumption is the code to be repaired contains no syntactic errors.

In this paper, we introduce MMAPR, a **multi-modal automated repair system**. MMAPR is a *unified* syntax and semantic repair engine for introductory Python programming assignments. We use a large language model trained on code (LLMC) as the core component in MMAPR. Using an LLMC removes the need for custom symbolic repair logic or retraining of a new neural model, and it allows us to handle *both* syntactic and semantic mistakes. While LLMCs have been successfully applied to tasks such as code generation [14], its impact in the education domain remains controversial [15]. Using an LLMC for repair provides an opportunity to produce a positive impact in this domain.

We follow the approach of recent work [16], [17] in framing program repair as a code generation task that can be tackled with an LLMC. However, using LLMCs to produce student repairs requires addressing three challenges. First, in the classroom setting, multiple sources can help to repair students’ code. For example, the instructor might provide test cases that the solution code needs to pass, and/or a description of the task in a natural language. Even the compiler messages can be used to repair syntax errors in the code. While standard LLMC-based APR tools take as input the buggy code and

This work was performed when Jialu Zhang was an intern at Microsoft Prose Team. All other authors are listed in alphabetical order.

only one correctness criterion (for example, test cases), our approach is multi-modal. In addition to the student’s buggy program, we also take all the above-mentioned sources as a part of the input. Moreover, we make effective use of other students’ submissions, if they are available. Second, we need to mitigate the extent to which the LLMC can generate more code than necessary or make changes to parts of the program that are not incorrect, which could result in an excessive patch. Third, using an LLMC as a black box means that we also need to adapt traditional prompt engineering techniques. We use *prompt-based learning* [18], which consists of creating a text-based template to generate the code.

MMAPR ensembles multi-modal prompts to generate complementary repair candidates. It employs prompts in an iterative querying strategy that first uses syntax-targeted prompts and *then* semantics-targeted prompts. **MMAPR** takes inspiration from existing symbolic repair literature [9], [12], [19] and leverages few-shot learning, which adds task-related examples to the prompt, by retrieving other student’s programs that have similar mistakes (and eventual corrections). To identify these programs, **MMAPR** computes a similarity metric over test-suite outcomes. Finally, to reduce the number of changes induced by syntax errors that should have relatively simple fixes, **MMAPR** uses the program’s structure to extract a subprogram to give as input to the LLMC. By reducing the code surface exposed to the LLMC, **MMAPR** biases repairs towards fewer edits.

We evaluated **MMAPR** on student programs, from an introductory Python programming course at a major university in India. Our evaluation has 15 programming tasks, totalling 286 student programs. These student programs contain both syntactic and semantic mistakes. As there is currently no tool that can solve both errors simultaneously, we combined BIFI [20] with Refactory [11] to create a state-of-the-art baseline system. BIFI is a state-of-the-art transformer-based syntax repair engine for Python, while Refactory is a state-of-the-art symbolic semantic repair engine for introductory Python assignments.

Our results show that **MMAPR** can effectively repair student programs in our benchmark set. While the baseline repairs 67.13% of programs, **MMAPR** (without few-shot learning) can repair 86.71%. If we add few-shot learning to **MMAPR**, this repair rate rises to 96.5%. Furthermore, the average token edit distance associated with **MMAPR** patches are smaller (31.4 without few-shots and 31.29 with few-shots) compared to the patches produced by the baseline (42.50).

We carried out an ablation study to understand the impact of our design decisions. Our results indicate that by performing iterative querying the repair rate rises from 82.87% to 86.71%. Furthermore, adding few-shots raises the repair success rate to 96.5%. The evaluation also shows that our techniques are important for maintaining the repaired program similar to the buggy input program. For example, removing the program chunker, which selects subprograms in the syntax repair phase, raises the average token edit distance from 5.46 to 9.38. We also show that different multi-modal prompts have varying performance, but if we combine their candidates as we do in **MMAPR**, we obtain the best performance.

To summarize, we make the following contributions:

- We propose an approach to automatically repair mistakes in students’ Python programming assignments using a large language model trained on code (LLMC). Our approach uses multimodal prompts, iterative querying, test-case-based few-shot selection, and structure-based program chunking to repair student mistakes. In contrast to prior work, our approach uses the *same* underlying LLMC to repair *both* syntactic and semantic mistakes.
- We implement this approach in **MMAPR**, which uses OpenAI’s popular Codex as the LLMC. We evaluate **MMAPR** on a dataset of 286 real student Python programs drawn from an introductory Python programming course in India. We compare performance to a baseline produced by combining a state-of-the-art syntax repair engine (BIFI) and a state-of-the-art semantic repair engine (Refactory). Our results show that **MMAPR** can outperform our baseline, even without few-shot learning. In addition, on average, **MMAPR**’s patches are closer to the original student submission.

The remainder of the paper is structured as follows. Section II walks through multiple examples of real student mistakes, as well as the associated **MMAPR** patches. Section III provides a brief background on concepts related to large language models. Section IV describes our approach in detail and its implementation in **MMAPR**. Section V provides experimental results on our dataset of student Python programs, including a comparison to a baseline repair system. We discuss related work in Section VII. Finally, we conclude with takeaways in Section VIII.

II. MOTIVATING EXAMPLE

Consider Figure 1, which shows a student’s incorrect program, along with a solution generated by **MMAPR**. The student is solving the task of reading two numbers from stdin and printing different things depending on whether both, either, or neither of the values are prime.

The student has made both syntactic and semantic mistakes. Lines 1 and 2 call `input` twice to read from stdin, and parse these values as integers using `int`. However, this constitutes a semantic mistake, as the assignment input format consists of two values *on the same line* separated by a comma. Furthermore, a traditional semantic repair engine would fail to fix this student’s assignment as there is *also* a syntactic mistake at line 30. The student used a single `=` for comparison in the `elif` clause (the correct syntax would be a double equals).

The **MMAPR** solution, shown alongside it, fixes the input processing (semantic mistake) by reading from stdin, splitting on the comma, and applying `int` (to parse as integer) using the `map` combinator. Line 23 fixes the syntax error by replacing single equals with double equals (for comparison). Interestingly, the underlying LLMC (Codex) *also* refactored the student’s program. In this case, lines 8 through 17 correspond to a function to check if a number is prime. This function is called twice, at lines 18 and 19. This replaces the repeated code in the original program, which spanned lines 9-17 and lines 17-26.

```

1 a=int(input())
2 b=int(input())
3 n=str(a)
4 m=str(b)
5 n1=n[::-1]
6 m1=m[::-1]
7 x=int(n1)
8 y=int(m1)
9 if x>1:
10     for i in range(2, x):
11         if (x%i)==0:
12             p=0
13             break
14         else:
15             p=1
16 else:
17     p=0
18 if y>1:
19     for j in range(2, y):
20         if (y%j)==0:
21             q=0
22             break
23         else:
24             q=1
25 else:
26     q=0
27 if p==1 and q==1:
28     g=x+y
29     print(g)
30 elif p=1 or q=1:
31     h=a+b
32     print(h)
33 else:
34     print(a*b)

```

(a) Incorrect Program

```

1 a,b = map(int,
2     input().split(',')
3     )
4 n=str(a)
5 m=str(b)
6 n1=n[::-1]
7 m1=m[::-1]
8 x=int(n1)
9 y=int(m1)
10 def prime(z):
11     if z>1:
12         for i in range(2,z):
13             if (z%i)==0:
14                 return 0
15             else:
16                 return 1
17         else:
18             return 0
19 p=prime(x)
20 q=prime(y)
21 if p==1 and q==1:
22     g=x+y
23     print(g)
24 elif p=1 or q==1:
25     h=a+b
26     print(h)
27 else:
28     print(a*b)

```

(b) MMAPR's Repair

Fig. 1: A student's submission contains both syntax and semantic mistakes (red). MMAPR's fixes (blue) the original semantic and syntactic issues and also refactors part of the student's code into a function (lines 8 - 17 in (b)) that avoids code duplication (lines 9-17, 18-26 in (a)).

The edit distance between the MMAPR repair and the original student program is 95, while the distance between the instructor's reference solution and the original student program is 188. A smaller edit distance is a key goal for APR in the educational domain, as this can help the student understand the repair with respect to their own mistakes.

Figure 2 presents another example of an incorrect student program and a solution generated by MMAPR. In this assignment, the students need to check whether a string, read from stdin, is a palindrome or not, and print out a message accordingly to stdout. For this student's program, MMAPR has to generate a complex repair that fixes four syntax mistakes and multiple semantic bugs.

The student has made syntax errors on lines 4, 8, 10, and 12, where they have left off the colon symbol necessary for control flow statements in Python. On line 2, the student called a non-existent function `lower`. The student has used standard division on lines 5, 6, 13, and 14 when they should have used integer division. The student has included two spurious print statements, at lines 7 and 15, which will interfere with

```

1 i = input()
2 S = i.lower()
3 l = len(S)
4 if(l%2!=0):
5     B = S[::(l+1)//2]
6     E = S[:l//2-1:-1]
7     #print(B,E)
8     if(B==E):
9         print(i, 'is a
10             palindrome.')
11     else:
12         print(i, 'is NOT a
13             palindrome.')
14 else:
15     B = S[:l//2]
16     E = S[:l//2-1:-1]
17     #print(B,E)
18     if(B==E):
19         print(i, 'is a
20             palindrome.')
21     else:
22         print(i, 'is NOT a
23             palindrome.')

```

(a) Incorrect Program

(b) MMAPR's Repair

Fig. 2: A complex repair (blue) that fixes multiple syntax and semantic mistakes (red). The repair produced by MMAPR, which preserves the overall structure of the student's program, makes fewer changes to the student's program than a patch with respect to the instructor's reference solution.

the instructor's test-suite execution, as the suite checks values printed to stdout for correctness. Finally, the student has omitted the expected print statements (along with the equality check) for the case where the input string is of even length.

While the student's program has many mistakes, the overall structure and key concepts are there. Looking at the MMAPR solution shown alongside, it resolves these mistakes but preserves the student's overall structure. In particular, MMAPR replaces the non-existent `lower` function with a call to the string method with the same name. It replaces the division operator (`/`) throughout the program with the intended floor division operator (`//`), comments out the extract print statements, and adds the missing equality check and print statements in the case of even-length inputs.

The edit distance between the MMAPR repair and the original student program is 52, while the distance between the instructor's reference solution and the original student program is 97. The reference solution is a standard one line program for palindrome. Once again, the MMAPR repair is closer to the student submission than the instructor's reference solution.

III. BACKGROUND

We now provide a short background on concepts related to large language models.

A *large language model* (LLM) can be viewed as a probability distribution over sequences of words. This distribution is learned using a deep neural network with a large number of parameters. These networks are typically trained on large amounts of text (or code) with objectives such as predicting

particular masked-out tokens or *autoregressive* objectives such as predicting the next token given the preceding tokens. When the LLM has been trained on significant amounts of code, we refer to it as a large language model trained on code (LLMC).

Often, LLMs are pre-trained and then *fine-tuned*, meaning trained further on more specialized data or tasks. A particularly popular LLMC is OpenAI’s Codex [21], a variant of GPT-3 [22] that is fine-tuned on code from more than 50 million GitHub repositories.

In contrast to traditional supervised machine learning, LLMs have shown to be effective for *few-* and even *zero-shot* learning. This means that the LLM can perform tasks it was not explicitly trained for just by giving it a few examples of the task or even no examples, respectively, at inference time.

In this setting of few- (or zero-)shot learning, the LLM is typically employed using what is termed *prompt-based learning* [18]. A *prompt* is a textual template that can be given as input to the LLM to obtain a sequence of iteratively predicted next tokens, called a *generation*. A prompt typically consists of a query and possibly zero or more examples of the task, called shots. For example, the prompt below includes a specific query to fix a syntax error. One valid generation, that fixes the syntax error, would be `print()`.

```
# Fix the syntax error of the program #
# Buggy program #
print("\")
```

In practice, a prompt can incorporate anything that can be captured in textual format. In particular, *multi-modal* prompts are those that incorporate different *modalities* of inputs, such as natural language, code, and data.

Different prompts may result in different LLM completions. Other factors may also affect the completions produced, such as the sampling strategy or hyperparameters for the sampling strategy. One important hyperparameter is *temperature*, which controls the extent to which we sample less likely completions.

While we use OpenAI’s Codex in this work, there are other such LLMs that could be used such as Salesforce’s CodeGen [23] or OpenScience’s BLOOM [24]. Even within OpenAI’s Codex there are different underlying models offered, including Codex-Edit [25]. We found performance to be better with the standard Codex completion model. We now leverage these concepts to describe our approach to APR.

IV. METHODOLOGY

Figure 3 provides an overview of the architecture underlying MMAPR. The student’s buggy program first enters a syntax repair phase. In this phase, we extract subprograms from the original program that have a syntax error. Each such subprogram is fed to a syntax prompt generator that produces multiple syntax-oriented prompts. The LLMC then generates repair candidates, which are validated by the syntax oracle. This process is repeated until all syntax errors are removed. Any candidate that has no syntax errors moves on to the semantic phase. In this phase, MMAPR uses a semantic prompt generator to produce semantics-oriented prompts. If it has access to other

student’s assignment history, MMAPR can also add few-shots to these prompts. These prompts are fed to the LLMC, which generates new program candidates. These are validated by the test-suite-based semantic oracle. If multiple candidates satisfy all test cases, MMAPR returns the one with the smallest token edit distance with respect to the student’s original buggy program. We now describe each step in detail.

A. Syntax Phase

Students typically first resolve syntax errors in their assignments, and then move on to resolve semantic errors (such as test case failures). MMAPR takes inspiration from this approach and similarly splits its repair into syntax and semantic phases.

In the first phase, MMAPR receives the student’s buggy program. A syntax oracle, for example, the underlying Python parser, is used to determine if there is a syntactic mistake. If there is no such mistake, the program can move into the semantic phase. However, if there is a mistake, MMAPR must produce a patch that resolves it, before moving to the semantic phase.

While our syntax prompt generator could directly include the original program in the prompt, we have found that doing so can result in spurious edits that are not actually necessary to resolve the syntax error. Existing work has also observed similar phenomena in the related area of natural language to code generation [26]. As a result, we introduced a component we call the *program chunker* to mitigate this challenge by reducing the amount of code included in the prompt.

1) *Program Chunking*: For each syntax mistake in the original buggy program, the program chunker extracts a subset of lines that contains (1) the oracle-reported syntax error location and (2) the nearest encompassing control-flow statement. These *chunks* are a heuristic approximation of a basic block, and allow us to restrict the code input given to the LLMC. Note that we perform this heuristic approximation as a standard analysis to extract basic blocks typically requires a syntactically correct input program.

Algorithm 1 Chunker: slicing the program snippet that contains the error message

Input: sC : Program Source Code

Input: msg : Compiler Message

Output: $chunkedCode$: Chunked Program Source Code

```
1: procedure CHUNKER( $sC, msg$ )
2:    $listsC, errorLine = locateError(sC, msg)$ 
3:    $indentLevel = getIndentationLevel(listsC, errorLine)$ 
4:   ▷ slice adjacent code with no less than the indentation level
5:    $start, end = sliceBiway(listsC, errorLine, indentLevel)$ 
6:   ▷ adjust slice pointers to account for containing control-flow
7:   if  $listsC[start].startswith(cfNodes)$  then
8:      $newindentLevel = getIndentationLevel(listsC, start)$ 
9:      $start, end = sliceBiway(listsC, start, newindentLevel)$ 
10:  return  $chunkedCode = listsC[start, end]$ 
```

MMAPR extracts the program chunk for the first (top-down) syntax error reported. Algorithm 1 outlines the procedure used to produce this program chunk. It takes advantage of

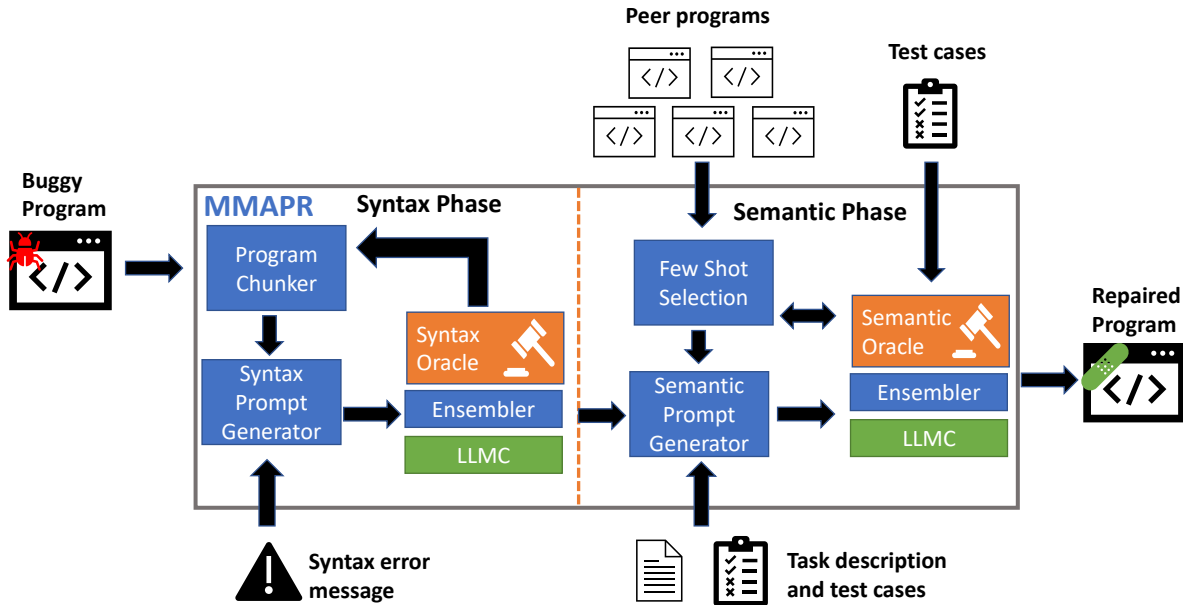


Fig. 3: MMAPR architecture. A buggy program first enters a syntax repair phase. In this phase, MMAPR transforms the program using a program chunker, which performs a structure-based subsetting of code lines to narrow the focus for the LLMC. Multiple syntax-oriented prompts are generated using this subprogram, fed to an LLMC, and any patches are integrated into the original program. If any candidate satisfies the syntax oracle, it can move on to the semantic phase. In the semantic phase, MMAPR leverages both the natural language description of the assignment and the instructor-provided test cases to create various prompts. In addition, if available, MMAPR can use other peers’ solutions as few-shots by selecting them using test-case-based selection to identify failures that resemble the current student’s program, along with eventually correct solutions. Prompts are fed to the LLMC to generate candidates. If multiple candidates satisfy the test suite, MMAPR returns the one with the smallest edit distance with respect to the original student program.

both control-flow structure (based on Python keywords) and indentation, which is meaningful in the Python language. The program chunker first identifies the adjacent code that has the same or large indentation level as the line with the syntax error. Then, if the code chunk contains control flow related keywords, such as `if` and `elif`, MMAPR makes sure the corresponding keywords (such as `elif` or `else`) is also in the chunk. This program chunk is then provided to the syntax prompt generator.

2) *Syntax Prompt Generator*: The syntax prompt generator produces two (multimodal) prompts, one with and one without the syntax error message reported by the syntax oracle. An example of both is shown in Figure 4. Because the syntax oracle is available, we do not need to choose a single prompt template for all programs, but instead we query the LLMC with both prompts, extract the code portion from each generation, merge it into the original program by replacing the lines corresponding to the current program chunk, and then rely on the syntax oracle to filter out invalid repairs.

If a program candidate has no syntax errors, it can move on to the semantic phase. If any syntax errors remain, the syntax phase is repeated on this candidate program. This iteration allows the repair of multiple, spatially-independent, syntax errors. For our evaluation, we allow this procedure to iterate at most two times to limit repair times.

| | |
|---|---|
| <pre># Buggy Program # while (n > 0): a = n % 10 ...</pre> | <pre>### Error Msg ### File "<unknown>", line 2 a = n % 10 ^ IndentationError: expected an indented block</pre> |
| (a) without error message | (b) with error message |

Fig. 4: The syntax prompt generator produces prompts that can include the buggy program or the error message. We elide portions of the code fragments for brevity.

B. Semantic Phase

After MMAPR has generated syntactically valid candidate programs, the repair procedure moves to a semantic repair phase. Intuitively, this phase incorporates information that allows the LLMC to generate candidate programs that satisfy the programming assignment task, as determined by a semantic oracle. Following the approach of existing work in automated repair for programming assignments [9], [11], we use the

```

[[Buggy Program]]
### Buggy Program ###
x=input()
y=int(x)
z = number % 10
y = 10 * y + z
number = number / 10
number = int(number)
print("Reverse: {}".format(x[::-1]))
print("Sum: {}".format(Sum))

[[Problem Description]]
#Write a program to read a number (int) from the
  user. Print the number in reverse. Also print
  the sum of the number and its reverse in a
  separate line. See the examples.
#NOTE: Do not print any prompt in the input().

[[Test Suite]]
#input:
43
#output:
Reverse: 34
Sum: 77

#input:
500
#output:
Reverse: 5
Sum: 505

### Correct Program ###

```

Fig. 5: An example multimodal prompt (in zero-shot setting for brevity) produced by the semantic prompt generator. This prompt includes code, natural language, and test cases. Lines starting with the double brackets are shown only for clarity.

instructor’s test suite (consisting of inputs and expected outputs) as the semantic oracle. We say a program has been repaired if it produces the expected outputs for the given inputs.

1) *Semantic Prompt Generator*: The semantic prompt generator takes advantage of the rich set of signals available in the education domain. In particular, we exploit the fact that programming assignments typically have available: (1) a natural language description of the task, (2) a set of test cases, and (3) peers’ programming solutions.

The semantic prompt generator takes as input a syntactically valid program, the task description in natural language, and the set of instructor-provided test cases. The generator then produces prompts with different combinations of this information. Figure 5 shows an example of such a multimodal prompt. This prompt includes the student’s buggy code, the natural language description of the assignment, as well as the input-output-based test cases.

If MMAPR has access to other students’ assignment solution history, then it can also employ few-shot learning, described in the following Section IV-B2, in each of these prompts.

Similarly to the syntax phase, rather than picking a single prompt template, we use all prompts generated and rely on the semantic oracle to identify viable repair candidates. Each prompt given to the LLMC can generate up to K candidates,

where we heuristically set K to ten to balance exploration of candidates with search space explosion. Each of these candidates is given to the semantic oracle, which executes that candidate on the test suite. We remove any candidate programs that result in a runtime exception or fail to satisfy any test cases.

If there are multiple valid candidate programs after the semantic phase, we return the one with the smallest token-based edit distance [20] to the student as the repaired program.

2) *Few-Shot Learning*: If MMAPR has access to other students’ programs it can employ few-shot learning. In contrast to other repair systems, such as Refactory [11], that typically employ only correct programs, MMAPR’s few-shots consist of both correct *and* incorrect programs.

In particular, MMAPR’s few-shot learning example bank consists of pairs of program versions (p, p') where both p and p' satisfy the syntax oracle, p' satisfies the semantic oracle but p does not, and p is a historical edit-version ancestor of p' . Given a candidate program produced by the syntax phase of MMAPR, we retrieve the three most similar p and their associated correct versions p' to include as shots in the LLMC prompts produced by the semantic prompt generator.

We take inspiration from traditional automated program repair and say two programs are similar if they result in similar test suite executions. We define a test suite execution vector for program p that captures test failures as $T_p \in \mathbb{B}^n = (t_0, \dots, t_n)$ where n is the number of test cases, and t_i is the boolean failure status of the i th test. We define the similarity function between p_1 and p_2 as $1 - \text{HAMMING}(T_{p_1}, T_{p_2})$ with HAMMING the normalized Hamming distance [27] between the two vectors.

Figure 6 shows an illustrative buggy student program, along with a single shot, consisting of the incorrect and corrected program from a peer. The shot is picked as it has the same test execution vector as our target buggy program.

Note that if MMAPR does not have access to peer programs, then it can still query the LLMC using a zero-shot approach. In our evaluation (Section V) we show that this ablated strategy still performs competitively.

V. EVALUATION

We explore the following two research questions in our evaluation of MMAPR:

- **(RQ1)** How does MMAPR’s overall repair rate (syntax and semantics) compare to a state-of-the-art baseline?
- **(RQ2)** What is the impact of the underlying design decisions in MMAPR? Specifically, what is the impact of the structure-based program chunking, iterative querying, test-case-based few-shot selection, and multimodal ensembled prompts?

Implementation. We have built a MMAPR prototype using a mix of Python and open-source software libraries. The core of MMAPR’s implementation consists of approximately 600 lines of Python code, which is 5 to 10 times less than a typical symbolic repair system in the education domain [9]–[11]. In addition to the reduced engineering efforts, MMAPR can handle

```

[[Shot Starts]]
# Incorrect Program #
print (m+n)
# Correct Program #
print (m*n)
[Shot Ends]

[[Buggy Program Starts]]
### Buggy Program ###
sum = m
i = 0
while i < n:
    sum += 1
    i += 1
print (sum)
[[Buggy Program Ends]]

[[Test Suite Starts]]
#input:
2 2
#output:
4

#input:
2 3
#output:
6
[[Test Suite Ends]]

### Correct Program ###

```

Fig. 6: An illustrative example of few-shot learning in MMAPR. Both the incorrect program in the shot and the target buggy program to be repaired have the same test suite execution results [pass, fail].

both syntax and semantic bugs in one system, while most systems focus on one of the two bug classes.

We selected the top 10 program candidates in each syntax and semantics phase based on the on the average token log probabilities produced by the LLMC.

For the model selection, we used OpenAI’s Codex as our LLMC. Specifically, we used the completion model. We found that other models, such as Codex Edit [25], did not perform as well. We set the temperature to 0.8 based on preliminary experiments.

Benchmark. We derived a benchmark set by selecting programs from a collection of introductory Python assignments collected by third-party authors in a large Indian university. This dataset is a Python-version of the dataset described in [28].

The dataset contains 18 assignments, each with a problem description, the test suite, and students’ authoring history. A student’s history consists of an ordered collection of program versions, where each version can be an explicit submission to the testing server, or a periodic (passive) snapshot – the dataset does not have a way to distinguish between these.

We removed three assignments that required reading files that are not reported in the dataset or that asked students to generate a PDF plot, which makes assessing correctness difficult without manual inspection. For each assignment, we selected the students that had an eventually correct program. For each such student, we collected the latest (closest to the

correct version in time) version that had a syntactic mistake. If the student never made a syntactic mistake, we remove this student from our benchmark. This results in a total of 286 *program pairs*, consisting of a buggy and ground-truth correct program version.

Baseline. Most repair systems focus on either syntax repairs, or semantic repairs.¹ To create a state-of-the-art baseline that performs *both* syntax and semantic repairs, we combined BIFI, a state-of-the-art transformer-based Python syntax repair tool, and Refactory, a state-of-the-art semantics repair tool designed for introductory Python assignments.

To run this baseline, we gave BIFI the original student program with syntax errors and generated 50 candidate programs for each buggy program. For each candidate, we ran the syntax oracle and checked for syntactic correctness. For each candidate that passed the syntax check, we called Refactory along with the instructor’s reference solution.² If Refactory can repair *any* of the candidates, we say it has repaired the student’s program. If there are multiple candidate programs that passed the test suite, we choose the one with the smallest token edit distance with respect to the buggy program as the final repair. We ran all experiments on a Windows VM with an Intel i7 CPU and 32 GB of RAM.

A. RQ1: Overall Repair Performance

Table I shows that without few-shot learning MMAPR can repair 86.71% of student programs. In contrast, our baseline repairs 67.13% of student programs. In addition, MMAPR repairs are closer to the original student programs. In particular, the mean token edit distance between the buggy program and our repaired program is 31.40 compared to 42.50 for our baseline.

Our results also show that if we add few-shot learning, by leveraging the availability of other student’s incorrect and correct versions paired with our test-case-based few-shot selection strategy, we can raise MMAPR’s overall repair rate to 96.50%. It is worth noting that the increase in repairs does not raise average token edit distance (31.29). In this setting, MMAPR outperformed the baseline by 29.37 percentage points in terms of repair rates and reduced average token edit distance by 26.38%.

A key idea behind Refactory is to perform a control-flow match between correct programs and the student’s buggy program. To do this, Refactory generates multiple versions of the correct program (in our case, the reference solution) by applying rewrites. Repairs that require substantial refactoring, such as that shown in the motivating example (Section II) are beyond Refactory’s expressiveness.

Repairing semantic mistakes typically depends on first resolving any syntactic mistakes. Indeed, students often focus on resolving mistakes reported by the parser/compiler before they move on to debugging test-case failures. MMAPR’s architecture,

¹A notable exception in the education domain is sk_p [8], however this tool is not publicly available and the repair rate (29%) described in the paper is low compared to our baselines.

²The original Refactory paper shows that there is little-to-no performance difference between providing one and multiple correct reference programs.

TABLE I: MMAPR (without few shots) repairs a larger fraction of programs (86.71%) compared to our baseline (67.13%). On average, MMAPR repairs are closer in terms of token edit distance (TED) to the original student program (31.40 versus 42.50). Adding few-shots based on other peers’ programs raises MMAPR’s repair rate to 96.50% while keeping a comparable average token edit distance (31.29).

| Method | | MMAPR (without few-shot learning) | | MMAPR (with few-shot learning) | | BIFI + Refactory | |
|------------|---------------|-----------------------------------|----------------|--------------------------------|----------------|------------------|----------------|
| Problem ID | # Submissions | Repair rate (%) | Mean TED (SD) | Repair rate (%) | Mean TED (SD) | Repair rate (%) | Mean TED (SD) |
| 2865 | 11 | 100.00 | 6.45 (4.74) | 100.00 | 6.45 (4.74) | 100.00 | 5.28 (4.27) |
| 2868 | 28 | 85.71 | 8.79 (8.94) | 100.00 | 8.64 (8.49) | 82.14 | 8.35 (7.00) |
| 2869 | 23 | 95.65 | 16.68 (18.47) | 100.00 | 10.30 (10.99) | 69.57 | 17.81 (13.67) |
| 2870 | 27 | 74.07 | 10.00 (13.33) | 100.00 | 15.00 (19.35) | 85.19 | 15.74 (23.92) |
| 2872 | 18 | 100.00 | 8.33 (15.15) | 100.00 | 7.39 (13.01) | 72.22 | 76.61 (9.24) |
| 2873 | 32 | 78.13 | 12.00 (16.18) | 90.63 | 12.93 (15.47) | 84.38 | 58.48 (14.04) |
| 2874 | 16 | 100.00 | 9.56 (12.50) | 100.00 | 8.50 (11.76) | 87.50 | 23.43 (16.28) |
| 2875 | 23 | 86.96 | 14.75 (19.97) | 100.00 | 11.52 (12.52) | 78.26 | 35.72 (14.58) |
| 2877 | 21 | 100.00 | 9.71 (16.82) | 100.00 | 9.14 (16.79) | 80.95 | 33.12 (24.56) |
| 2878 | 25 | 100.00 | 37.00 (60.16) | 100.00 | 36.32 (59.53) | 68.00 | 97.94 (23.79) |
| 2879 | 21 | 76.19 | 131.19 (51.62) | 85.71 | 132.78 (52.61) | 52.38 | 125.64 (11.33) |
| 2882 | 23 | 60.87 | 90.64 (71.76) | 91.30 | 106.57 (77.57) | 0.00 | N/A |
| 2883 | 5 | 100.00 | 17.40 (14.67) | 100.00 | 17.40 (14.67) | 40.00 | 53 (0.00) |
| 2920 | 10 | 80.00 | 84.38 (67.62) | 80.00 | 53.50 (66.05) | 0.00 | N/A |
| 2921 | 3 | 100.00 | 28.00 (3.61) | 100.00 | 28.00 (3.61) | 0.00 | N/A |
| Overall | | 86.71 | 31.40 | 96.50 | 31.29 | 67.13 | 42.50 |

TABLE II: The first stage in the repair process is to fix syntax errors. MMAPR can produce a syntactically valid candidate for all programs in our benchmark, compared to 80.07% for BIFI. On average, MMAPR’s repairs are also closer to the original program (edit distance of 5.46 versus 25.07).

| Method | | MMAPR | | BIFI | |
|---------|-------|---------------|-----------------|----------------|-----------------|
| ID | # Sub | Mean TED (SD) | Repair rate (%) | Mean TED (SD) | Repair rate (%) |
| 2865 | 11 | 2.18 (1.25) | 100.00 | 1.82 (0.75) | 100.00 |
| 2868 | 28 | 2.75 (2.17) | 82.14 | 1.83 (1.11) | 82.14 |
| 2869 | 23 | 2.91 (2.41) | 73.91 | 1.47 (0.80) | 73.91 |
| 2870 | 27 | 2.33 (2.18) | 85.19 | 2.04 (1.89) | 85.19 |
| 2872 | 18 | 2.39 (1.2) | 72.22 | 1.62 (0.87) | 72.22 |
| 2873 | 32 | 2.84 (2.58) | 84.38 | 2.56 (2.04) | 84.38 |
| 2874 | 16 | 2.06 (1.84) | 87.50 | 2.07 (2.02) | 87.50 |
| 2875 | 23 | 2.78 (2.71) | 78.26 | 1.78 (1.56) | 78.26 |
| 2877 | 21 | 2.19 (1.29) | 80.95 | 3.18 (7.47) | 80.95 |
| 2878 | 25 | 4.84 (8.58) | 0.00 | 40.2 (60.65) | 0.00 |
| 2879 | 21 | 18.86 (21.24) | 66.67 | 117.00 (58.15) | 66.67 |
| 2882 | 23 | 17.39 (23.23) | 86.96 | 127.65 (74.76) | 86.96 |
| 2883 | 5 | 5.60 (9.74) | 80.00 | 36.25 (37.98) | 80.00 |
| 2920 | 10 | 10.30 (18.68) | 50.00 | 51.75 (54.90) | 50.00 |
| 2921 | 3 | 1.67 (0.58) | 100.00 | 1.33 (0.58) | 100.00 |
| Overall | | 5.46 | 80.07 | 25.07 | 80.07 |

as well as our composed baseline, reflects this approach. As a result, we also want to understand syntax repair performance.

Table II summarizes the *syntax repair rates* across assignments and approaches. Our results show that MMAPR repairs the syntax bugs in all of the 286 programs, with a 100% syntax repair rate. This outperforms the state-of-the-art BIFI, which has a syntax repair rate of 80.07%. In addition, MMAPR’s repairs have a substantially lower mean token edit distance (5.46 versus 25.07), meaning our repairs on average introduce fewer changes to the original programs, which may facilitate understanding of the fixes.

We also observed that in 17 out of 286 cases, BIFI fails to handle the input program, potentially due to lexer issues. This highlights another advantage of using MMAPR to repair programs because MMAPR does not have any constraints over the input as a result of its prompt-based learning strategy.

BIFI is very effective at repairing small syntax mistakes in assignments of lower difficulty. For example, in assignment 2865, BIFI repairs all syntax errors and does so with a smaller average token edit distance (1.82 versus 2.18) compared to MMAPR. One interesting direction for future work is to combine BIFI with MMAPR, as the repairs can be complementary. In this case, Codex could focus on generating more complex repairs and BIFI could focus on small edits for simpler tasks such as missing a quote in a string.

B. RQ2: Ablation Study

We now present the results of experiments to analyze different design choices in MMAPR. MMAPR uses multimodal prompts, iterative querying, test-case-based few-shot selection, and structure-based program chunking to repair student mistakes. The power of few-shot selection was already shown in Table I. We will now present the results of the other three design choices.

1) *Program Chunking*: In the syntax stage, MMAPR first extracts program chunks from the original buggy program as detailed in Section IV. The intuition is that these chunks contain the syntax error we want to fix, along with surrounding context, while excluding code lines that are not relevant to the fix. Our goal is to reduce the number of (spurious) edits produced by the LLMC by reducing the code surface in the prompt.

To evaluate the impact of program chunking, we removed it from MMAPR and compared performance to the original MMAPR. Table III shows the average token edit distance produced in

the syntax phase with and without program chunking. We found that program chunking can reduce the average token edit distance up to 56.32% (problem assignment 2878). Overall, the average token edit distance is reduced from 9.38 to 5.46 (41.79%) by adding program chunking.

TABLE III: Programs chunking reduces the average token edit distance across all assignments. PG is short for performance gain.

| Method | MMAPR (no chunking) | MMAPR (with chunking) | |
|---------|---------------------|-----------------------|--------|
| ID | Mean TED (SD) | Mean TED (SD) | PG (%) |
| 2865 | 2.45 (1.21) | 2.18 (1.25) | 11.11 |
| 2868 | 2.82 (2.14) | 2.75 (2.17) | 2.53 |
| 2869 | 2.91 (2.41) | 2.91 (2.41) | 0.00 |
| 2870 | 2.33 (2.18) | 2.33 (2.18) | 0.00 |
| 2872 | 2.44 (1.2) | 2.39 (1.2) | 2.27 |
| 2873 | 3.09 (2.61) | 2.84 (2.58) | 8.08 |
| 2874 | 2.25 (2.08) | 2.06 (1.84) | 8.33 |
| 2875 | 3.52 (4.13) | 2.78 (2.71) | 20.99 |
| 2877 | 2.29 (1.27) | 2.19 (1.29) | 4.17 |
| 2878 | 11.08 (20.3) | 4.84 (8.58) | 56.32 |
| 2879 | 33.14 (24.91) | 18.86 (21.24) | 43.10 |
| 2882 | 42.57 (41.54) | 17.39 (23.23) | 59.14 |
| 2883 | 6.20 (11.08) | 5.60 (9.74) | 9.68 |
| 2920 | 15.20 (19.45) | 10.30 (18.68) | 32.24 |
| 2921 | 1.67 (0.58) | 1.67 (0.58) | 0.00 |
| Overall | 9.38 | 5.46 | 41.79 |

2) *Iterative Querying*: Students typically resolve syntax errors first and then move on to resolving semantic mistakes. MMAPR’s architecture follows this same intuition. To compare the effectiveness of this iterative approach, we ran a variant of MMAPR that addresses both syntax and semantic bugs in a *single* round. Table IV shows the results of this ablated variant and full MMAPR (without few-shots). We find that splitting concerns into two phases results in an increase in the overall repair rate from 82.87% to 86.71%. Using two phases increases the average token edit distance slightly (30.29 to 31.40).

3) *Multimodal Prompts*: MMAPR combines different types of input (code, natural language, test cases) into its prompts. This richness of inputs is a particular advantage of the educational setting. MMAPR *ensembles* these various prompts by querying the LLMC and *then* relying on the (syntax or semantics) oracle to filter out candidates. This approach is based on the idea that different prompts may produce complementary candidates. To demonstrate this complementarity, we generate different prompt structures, use them in a single query to the LLMC, and evaluate the overall repair rate across assignments.

Figure 7 shows that different prompt structures result in different overall performances in terms of fix rate. If a single prompt structure needs to be chosen, Program + Diagnostics + Description + Tests structure is most effective in this experiment. However, if we *ensemble* the candidates, we obtain the best result as candidates are complementary.

4) *Combining BIFI and Refactory to Repair Programs*: We investigated why our baseline produces repairs that (on average) have a larger token edit distance compared to MMAPR. We found that in some cases, BIFI produces repairs by deleting

TABLE IV: MMAPR performs iterative querying, splitting the repair procedure into a syntax and a semantic phase. We find that this iterative approach raises the overall repair rate from 82.87% to 86.71% (without few-shots). RR stands for repair rate.

| Method | MMAPR (no iterative) | | MMAPR (with iterative) | |
|---------|----------------------|----------------|------------------------|----------------|
| ID | RR (%) | Mean TED (SD) | RR (%) | Mean TED (SD) |
| 2865 | 100.00 | 6.45 (4.74) | 100.00 | 6.45 (4.74) |
| 2868 | 85.71 | 8.92 (8.88) | 85.71 | 8.79 (8.94) |
| 2869 | 86.96 | 13.35 (12.36) | 95.65 | 16.68 (18.47) |
| 2870 | 70.37 | 11.42 (13.87) | 74.07 | 10.00 (13.33) |
| 2872 | 100.00 | 8.50 (15.22) | 100.00 | 8.33 (15.15) |
| 2873 | 71.88 | 9.48 (11.63) | 78.13 | 12.00 (16.18) |
| 2874 | 100.00 | 9.75 (12.51) | 100.00 | 9.56 (12.50) |
| 2875 | 82.61 | 13.16 (18.69) | 86.96 | 14.75 (19.97) |
| 2877 | 100.00 | 9.71 (16.82) | 100.00 | 9.71 (16.82) |
| 2878 | 100.00 | 38.16 (62.24) | 100.00 | 37.00 (60.16) |
| 2879 | 71.43 | 130.07 (53.23) | 76.19 | 131.19 (51.62) |
| 2882 | 56.52 | 97.85 (72.64) | 60.87 | 90.64 (71.76) |
| 2883 | 100.00 | 17.40 (14.67) | 100.00 | 17.40 (14.67) |
| 2920 | 50.00 | 50.20 (48.9) | 80.00 | 84.38 (67.62) |
| 2921 | 100.00 | 28.00 (3.61) | 100.00 | 28.00 (3.61) |
| Overall | 82.87 | 30.29 | 86.71 | 31.40 |




Fig. 7: Rather than pick a single prompt structure, MMAPR *ensembles* them by querying the LLMC with multiple prompts, and then relying on the (syntax and semantic) oracles to rule out invalid candidates. Ensembling outperforms any particular prompt structure as candidates are complementary.

a portion of the code snippet that contains the syntax errors. Although this is an effective way to deal with syntax errors, it makes repairing semantic errors harder by deleting parts that may capture crucial logic in the original source code.

Below is one such example from our evaluation. The code snippet contains a syntax mistake in the last line. The parser complains that the “*Expression cannot contain an assignment =*”. In particular, the student has written an equals when they should have used a plus operator (which corresponds to the repair produced by MMAPR).

```

1 marksSum={}
2 for i in total:
3     if int(i[0])not in marksSum:
4         marksSum[int(i[0])]= int(i[2])
5     else:
6         k=int(i[2])
7         marksSum[int(i[0])]+=k
8 for i in sorted(marksSum):
9     print(str(i)+" "+str(marksSum[i]))

```

However, BIFI produced a different fix by removing the second `for` loop (lines 8-9) completely. This deletion introduced challenges for Refactory in the later semantic repair phase. Although Refactory in the end successfully repaired this program, the repair it generated is syntactically equivalent to the reference solution and is effectively completely rewritten with respect to the original incorrect program.

Overall, our comparison between MMAPR and BIFI + Refactory highlights the challenges in combining the state-of-the-art syntax and semantics tools to repair incorrect introductory programming assignments. BIFI and Refactory each focus on their targets, syntax bugs repair and semantic bugs repair, respectively. Combining them does not guarantee the best overall repair (as in the example shown above). Additionally, combining BIFI and Refactory required non-trivial engineering efforts, further motivating the need for a unified approach that can handle both types of bugs.

VI. THREATS TO VALIDITY

MMAPR validates candidate repairs by comparing execution results on the test suite with the reference program given by instructors. Validating program correctness through tests is not as strong as formal verification. To the best of our knowledge, the use of tests as a proxy for correctness is standard in the educational domain [9], [29].

We carried out our evaluation on one particular set of 286 student program. The size of the dataset is on par with the state-of-the-art automated program repair techniques [6], [30], but increasing the size of the evaluation dataset may provide additional insights and presents an opportunity for future work. We only evaluated MMAPR on Python programs. However, our prompt-based approach can be applied to other languages.

VII. RELATED WORK

Automated Program Repair. The software engineering and programming languages community has a long history of developing tools for automatically repairing buggy programs. Existing approaches have applied a variety of technical ideas, including formal program logic [5], [31], [32], search-based techniques [33] like genetic programming [2], [34], [35], and machine learning [3], [4], [36]. A particularly popular approach to APR consists of generating many program candidates, typically derived by performing syntactic transformations of the original buggy program, and then validating these candidates using test suite as an oracle [2].

Similarly, MMAPR aims to automatically fix syntax and semantic errors in buggy programs, and uses a syntax oracle (the compiler) and semantic oracle (test cases) to validate candidate programs produced. However, in contrast to existing work, MMAPR employs a large language model (Codex) as the main program transformation module and uses an ensemble of multi-modal prompts to improve its success rate. Additionally, students struggle even with basic syntax errors [37], [38]. Therefore, MMAPR targets students' incorrect submissions, rather than professional developers' production bugs. As a result, MMAPR has an additional requirement of minimizing the size

of the change made to allow students to better learn from the repaired program.

AI for Programming Education. AI, both symbolic and neural, has been extensively applied to the domain of education [39], [40]. In particular, for programming education past research has explored topics including feedback generation [9]–[12], [29], [41], [42] and program repair [20], [36], [43]–[46]. MMAPR is complementary to the series of work, showing that the task of program repair in this domain can be successfully tackled using an LLMC. Using such an approach can lower the effort to develop and maintain an APR tool for introductory programming assignments.

LLMs for Code Intelligence. Large pre-trained language models, such as OpenAI's Codex, Salesforce CodeGen [23], and BigScience's BLOOM [24], have been shown to be effective for a range of code intelligence tasks. For example, Microsoft's Copilot [14] builds on Codex to produce more effective single-line and multi-line code completion suggestions. Prior work has shown that such LLMs can also be used for repairing programs outside of the educational context [47]–[50]. Using these models to perform code generation from informal specifications, such as natural language, has also been a topic of active research [39]. Similarly to this work, MMAPR uses Codex to perform a code intelligence task (program repair). However, MMAPR is designed to focus on student programming tasks and as such our design decisions (e.g., strategies to reduce token edit distance) may not apply to other domains such as professional developers.

VIII. CONCLUSION

In this work, we introduced an approach to repair syntactic and semantic mistakes in introductory Python assignments. At the core of our approach sits a large language model trained on code. We leverage multimodal prompts, iterative querying, test-case-based few-shot selection, and program chunking to produce repairs. We implement our approach using Codex in a system called MMAPR and evaluate it on real student programs. Our results show that our unified system MMAPR outperforms a baseline, a combination of a state-of-the-art syntax repair engine and a state-of-the-art semantic repair engine, while producing smaller patches.

REFERENCES

- [1] N. Singer, "The hard part of computer science? getting into class," *The New York Times*, Jan. 2019, accessed: 2012-09-01.
- [2] C. Le Goues, T. Nguyen, S. Forrest, and W. Weimer, "Genprog: A generic method for automatic software repair," *IEEE Transactions on Software Engineering*, vol. 38, no. 1, pp. 54–72, 2012.
- [3] F. Long, P. Amidon, and M. Rinard, "Automatic inference of code transforms for patch generation," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 727–739.
- [4] F. Long and M. Rinard, "Automatic patch generation by learning correct code," in *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, ser. POPL '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 298–312.

- [5] S. Mechtaev, J. Yi, and A. Roychoudhury, "Angelix: Scalable multiline program patch synthesis via symbolic analysis," in *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*, 2016, pp. 691–701.
- [6] U. Z. Ahmed, Z. Fan, J. Yi, O. I. Al-Bataineh, and A. Roychoudhury, "Verifix: Verified repair of programming assignments," *ACM Trans. Softw. Eng. Methodol.*, vol. 31, no. 4, jul 2022.
- [7] C. Le Goues, M. Pradel, and A. Roychoudhury, "Automated program repair," *Commun. ACM*, vol. 62, no. 12, p. 56–65, 2019.
- [8] Y. Pu, K. Narasimhan, A. Solar-Lezama, and R. Barzilay, "Sk_p: A neural program corrector for moons," in *Companion Proceedings of the 2016 ACM SIGPLAN International Conference on Systems, Programming, Languages and Applications: Software for Humanity*, ser. SPLASH Companion 2016. New York, NY, USA: Association for Computing Machinery, 2016, p. 39–40.
- [9] S. Gulwani, I. Radicek, and F. Zuleger, "Automated clustering and program repair for introductory programming assignments," *SIGPLAN Not.*, vol. 53, no. 4, p. 465–480, Jun. 2018.
- [10] R. Rolim, G. Soares, L. D'Antoni, O. Polozov, S. Gulwani, R. Gheyi, R. Suzuki, and B. Hartmann, "Learning syntactic program transformations from examples," in *ICSE 2017*, April 2017.
- [11] Y. Hu, U. Z. Ahmed, S. Mechtaev, B. Leong, and A. Roychoudhury, "Refactoring based program repair applied to programming assignments," in *2019 34th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2019, pp. 388–398.
- [12] K. Wang, R. Singh, and Z. Su, "Search, align, and repair: Data-driven feedback generation for introductory programming exercises," ser. PLDI 2018. New York, NY, USA: Association for Computing Machinery, 2018, p. 481–495.
- [13] U. Z. Ahmed, P. Kumar, A. Karkare, P. Kar, and S. Gulwani, "Compilation error repair: for the student programs, from the student programs," in *Proceedings of the 40th International Conference on Software Engineering: Software Engineering Education and Training*, 2018, pp. 78–87.
- [14] "Microsoft Copilot." <https://github.com/features/copilot/>, Aug. 2022.
- [15] E. Berger, "Coping with copilot," <https://www.sigarch.org/coping-with-copilot/>, 2022.
- [16] C. S. Xia and L. Zhang, "Less training, more repairing please: Revisiting automated program repair via zero-shot learning," *arXiv preprint arXiv:2207.08281*, 2022.
- [17] H. Joshi, J. Cambroner, S. Gulwani, V. Le, I. Radicek, and G. Verbruggen, "Repair is nearly generation: Multilingual program repair with llms," *arXiv preprint arXiv:2208.11640*, 2022.
- [18] P. Liu, W. Yuan, J. Fu, Z. Jiang, H. Hayashi, and G. Neubig, "Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing," *arXiv preprint arXiv:2107.13586*, 2021.
- [19] Y. Ke, K. T. Stolee, C. Le Goues, and Y. Brun, "Repairing programs with semantic code search (t)," in *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 2015, pp. 295–306.
- [20] M. Yasunaga and P. Liang, "Break-it-fix-it: Unsupervised learning for program repair," in *ICML*, ser. Proceedings of Machine Learning Research, vol. 139. PMLR, 2021, pp. 11 941–11 952.
- [21] M. Chen *et al.*, "Evaluating large language models trained on code," 2021.
- [22] T. B. Brown *et al.*, "Language models are few-shot learners," in *NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [23] E. Nijkamp, B. Pang, H. Hayashi, L. Tu, H. Wang, Y. Zhou, S. Savarese, and C. Xiong, "A conversational paradigm for program synthesis," 2022. [Online]. Available: <https://arxiv.org/abs/2203.13474>
- [24] H. Laurençon, L. Saulnier, T. Wang, C. Akiki, A. V. del Moral, T. Le Scao, L. Von Werra, C. Mou, E. G. Ponferrada, H. Nguyen *et al.*, "The bigscience corpus a 1.6 tb composite multilingual dataset."
- [25] "New GPT-3 Capabilities: Edit & Insert," Mar. 2022. [Online]. Available: <https://openai.com/blog/gpt-3-edit-insert/>
- [26] G. Poesia, O. Polozov, V. Le, A. Tiwari, G. Soares, C. Meek, and S. Gulwani, "Synchromesh: Reliable code generation from pre-trained language models," *arXiv preprint arXiv:2201.11227*, 2022.
- [27] R. W. Hamming, "Error detecting and error correcting codes," *The Bell System Technical Journal*, vol. 29, no. 2, pp. 147–160, 1950.
- [28] D. Chhatbar, U. Z. Ahmed, and P. Kar, "Macer: A modular framework for accelerated compilation error repair," in *Artificial Intelligence in Education: 21st International Conference, AIED 2020, Ifrane, Morocco, July 6–10, 2020, Proceedings, Part I*. Berlin, Heidelberg: Springer-Verlag, 2020, p. 106–117.
- [29] R. Singh, S. Gulwani, and A. Solar-Lezama, "Automated feedback generation for introductory programming assignments," in *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 13. New York, NY, USA: Association for Computing Machinery, 2013, p. 15–26.
- [30] L. Li, H. Liu, K. Li, Y. Jiang, and R. Sun, "Generating concise patches for newly released programming assignments," *IEEE Transactions on Software Engineering*, pp. 1–1, 2022.
- [31] S. Mechtaev, M. Nguyen, Y. Noller, L. Grunske, and A. Roychoudhury, "Semantic program repair using a reference implementation," in *Proceedings of the 40th International Conference on Software Engineering, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, M. Chaudron, I. Crnkovic, M. Chechik, and M. Harman, Eds. ACM, 2018, pp. 129–139.
- [32] R. Shariffdeen, Y. Noller, L. Grunske, and A. Roychoudhury, "Concolic program repair," in *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation*, ser. PLDI 2021. New York, NY, USA: Association for Computing Machinery, 2021, p. 390–405.
- [33] C. Wong, P. Santiesteban, C. Kästner, and C. Le Goues, "Varfix: balancing edit expressiveness and search effectiveness in automated program repair," in *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ACM, 2021, pp. 354–366.
- [34] Y. Qi, X. Mao, Y. Lei, Z. Dai, and C. Wang, "The strength of random search on automated program repair," in *Proceedings of the 36th International Conference on Software Engineering*, ser. ICSE 2014, New York, NY, USA, 2014, p. 254–265.
- [35] D. Kim, J. Nam, J. Song, and S. Kim, "Automatic patch generation learned from human-written patches," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. IEEE Press, 2013, p. 802–811.
- [36] K. Wang, Z. Su, and R. Singh, "Dynamic neural program embeddings for program repair," in *International Conference on Learning Representations*, 2018.
- [37] A. Altadmri and N. C. Brown, "37 million compilations: Investigating novice programming mistakes in large-scale student data," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, ser. SIGCSE '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 522–527.
- [38] I. Drosos, P. J. Guo, and C. Parnin, "Happyface: Identifying and predicting frustrating obstacles for learning programming at scale," in *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, 2017, pp. 171–179.
- [39] Y. Li *et al.*, "Competition-level code generation with alphacode," Feb 2022.
- [40] J. Finnie-Ansley, P. Denny, B. A. Becker, A. Luxton-Reilly, and J. Prather, "The robots are coming: Exploring the implications of openai codex on introductory programming," in *Australasian Computing Education Conference*, ser. ACE '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 10–19.
- [41] D. Song, W. Lee, and H. Oh, *Context-Aware and Data-Driven Feedback Generation for Programming Assignments*. New York, NY, USA: Association for Computing Machinery, 2021, p. 328–340.
- [42] J. Zhang, D. Li, J. C. Kolesar, H. Shi, and R. Piskac, "Automated feedback generation for competition-level code," *CoRR*, vol. abs/2206.01848, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2206.01848>
- [43] J. Yi, U. Z. Ahmed, A. Karkare, S. H. Tan, and A. Roychoudhury, "A feasibility study of using automated program repair for introductory programming assignments," in *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering*, ser. ESEC/FSE 2017. New York, NY, USA: Association for Computing Machinery, 2017, p. 740–751.
- [44] Q. Xin and S. P. Reiss, "Leveraging syntax-related code for automated program repair," in *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE 2017. IEEE Press, 2017, p. 660–670.
- [45] E. Dinella, H. Dai, Z. Li, M. Naik, L. Song, and K. Wang, "Hoppity: Learning graph transformations to detect and fix bugs in programs," in *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.

- [46] Y. Lu, N. Meng, and W. Li, "FAPR: fast and accurate program repair for introductory programming courses," *CoRR*, vol. abs/2107.06550, 2021. [Online]. Available: <https://arxiv.org/abs/2107.06550>
- [47] G. Verbruggen, V. Le, and S. Gulwani, "Semantic programming by example with pre-trained models," *Proc. ACM Program. Lang.*, vol. 5, no. OOPSLA, pp. 1–25, 2021.
- [48] K. Rahmani, M. Raza, S. Gulwani, V. Le, D. Morris, A. Radhakrishna, G. Soares, and A. Tiwari, "Multi-modal program inference: a marriage of pre-trained language models and component-based synthesis," *Proc. ACM Program. Lang.*, vol. 5, no. OOPSLA, pp. 1–29, 2021.
- [49] E. Dinella, G. Ryan, T. Mytkowicz, and S. K. Lahiri, "TOGA: A neural method for test oracle generation," in *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*. ACM, 2022, pp. 2130–2141.
- [50] J. Zhang, T. Mytkowicz, M. Kaufman, R. Piskac, and S. K. Lahiri, "Using pre-trained language models to resolve textual and semantic merge conflicts (experience paper)," in *ISSTA*. ACM, 2022, pp. 77–88.