

Improving Code Performance Using LLMs in Zero-Shot: RAPGen

Spandan Garg*, Roshanak Zilouchian Moghaddam[†] and Neel Sundaresan[‡]

Microsoft Corporation

Email: *spgarg@microsoft.com, rozilouc@microsoft.com, [‡]neels@microsoft.com

Abstract—Performance bugs are non-functional bugs that can even manifest in well-tested commercial products. Fixing these performance bugs is an important yet challenging problem. In this work, we address this challenge and present a new approach called Retrieval-Augmented Prompt Generation (RAPGen). Given a code snippet with a performance issue, RAPGen first retrieves a prompt instruction from a pre-constructed knowledge-base of previous performance bug fixes and then generates a prompt using the retrieved instruction. It then uses this prompt on a Large Language Model in zero-shot to generate a fix. We compare our approach with the various prompt variations and state of the art methods in the task of performance bug fixing. Our empirical evaluation shows that RAPGen can generate performance improvement suggestions equivalent or better than a developer in $\sim 60\%$ of the cases, getting $\sim 42\%$ of them verbatim, in an expert-verified dataset of past performance changes made by C# developers. Furthermore, we conduct an in-the-wild evaluation to verify the model’s effectiveness in practice by suggesting fixes to developers in a large software company. So far, we have shared performance fixes on 10 codebases that represent production services running in the cloud and 7 of the fixes have been accepted by the developers and integrated into the code.

I. INTRODUCTION

Performance bugs are inefficient code snippets in software that can unnecessarily waste time and resources. Unlike functional bugs, performance bugs do not usually cause system failure. They tend to be harder to detect [Attariyan et al.(2012)], [Dean et al.(2014a)], [Dean et al.(2014b)], [Jovic et al.(2011)] and fix compared to functional bugs [Nistor et al.(2013)], [Song and Lu(2014)] and are usually fixed by expert developers [Chen et al.(2019)] and often manifest through large inputs or specific execution configurations [Nistor et al.(2013)], [Olivo et al.(2015)]. Therefore, they can sometimes go undetected for a long periods of time [Dean et al.(2014b)], [Jovic et al.(2011)]. As a result, better tool support is needed to fix performance bugs, specially for novice developers. For such a tool to be applicable in practice, it should also support a wide-range of performance bugs. However, the existing fixing approaches usually target specific kinds of performance bugs, such as repeated computations [Della Toffola et al.(2015)], software misconfigurations [Iqbal et al.(2021)], loop inefficiencies [Nistor et al.(2013)]. The majority of these approaches are also rule-based analyzers, with high maintenance cost [Bielik et al.(2017)]. Recently, to address this challenge, Garg et al. developed a deep-learning based performance fix tool called DeepDev-PERF [Garg et al.(2022)] that supports a

variety of performance bugs and does not require maintaining explicit rules or an expert system. Our approach builds on that work by showing the cost of building and maintaining such a tool can be further improved through the use of prompt engineering as opposed to intensive fine-tuning.

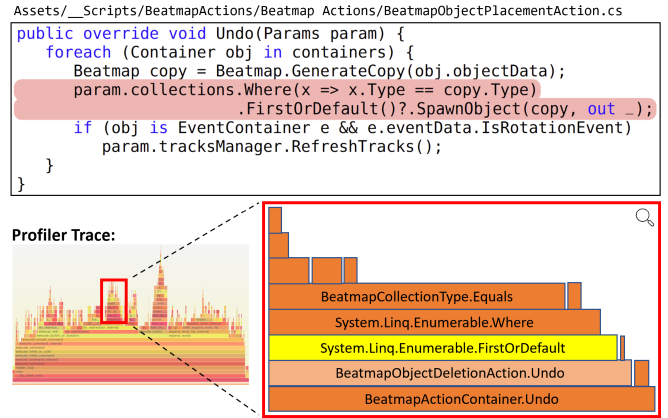


Fig. 1: A C# code snippet with an expensive LINQ query (highlighted in red) from a performance bug fix commit on GitHub. LINQ tends to get misused by developers and can often lead to performance issues, such as the one above. In this case, the LINQ methods `Where` and `FirstOrDefault` are used to iterate over a collection to find all entries matching a predicate, when only the first match is needed and the search could potentially stop early. Depending on the size of the collection and how frequently this code gets invoked, this may become a performance hot-spot in an application. The screenshot below shows flamegraph corresponding to this application, with the relevant section highlighted. The call-stack shows the `FirstOrDefault` (its corresponding frame being highlighted in yellow) as being the most expensive line within the `Undo` method.

With the rising adoption of tools such as Github copilot [Chen et al.(2021)] by developers, large language models are showing a promising success as productivity tools. Additionally, a promising paradigm for few-shot and zero-shot learning has emerged where a pre-trained large language model is adapted to various downstream tasks by conditioning it on natural language prompts. This paradigm has also been applied to a variety of tasks in Software Engineering, including the task of bug-fixing [Pearce et al.(2021b)], [Xia and Zhang(2022)].

Building upon this prior work, we present an approach called RAPGen that uses large language models (LLMs) in zero-shot, for the task of generating fixes for performance bugs using prompt engineering. Given a line of code that contains a performance bug (i.e. the buggy line), we compare this line with a pre-constructed knowledge-base to retrieve a prompt instruction that can be used to convey what change needs to be made to a LLM. The index contains patterns of buggy API usages extracted from a data-set of past performance fixes, along with an instruction explaining the fix transformation. Given a method with a performance issue and the corresponding buggy line within the method, we first retrieve an instruction from the knowledge-base and then construct an input prompt for the model, containing the buggy method, followed by the instructive comment and the method signature, coaxing the model to output the fixed version of the method. This type of zero-shot bug-fixing experience is desirable because (1) it avoids an expensive fine-tuning step, (2) alleviates the need to find a large high-quality labeled data-set of performance fixes for training such a model, and (3) fixes a wide-range of performance problems and could easily be extended to more bugs (by growing the number of commits used to build the knowledge-base) or other languages (by collecting relevant commits from other languages), without needing to train a separate model.

Evaluation and Effectiveness. We evaluated our approach on the most recent performance bug fix data set used in the DeepDev-PERF study [Garg et al.(2022)]. Experimental results show that our approach is indeed able to generate correct suggestions for 60% of the dataset, often within the very first suggestion, and out-performs DeepDev-PERF and other baselines for single method performance fixes. While designing an optimal prompt is still an open problem and remains elusive as ever, we also compare our approach with various prompt variations to demonstrate our design choices and our experiments with recent ideas such as reasoning-based prompting [Wei et al.(2022)]. Furthermore, through our "In-the-wild" evaluation, we show real-world evidence that RAPGen generates suggestions that lead to actual performance improvements to real .NET service codebases within a large company that are currently in production. These changes were validated and adopted by real developers.

II. MOTIVATING EXAMPLE

In this section, we provide an overview of RAPGen using a motivating example. Figure 1 shows a C# code method with a performance bug in the form of an expensive LINQ [Pialorsi and Russo(2007)] query (highlighted in red). The query iterates over a collection, finding all entries that match the predicate passed to the LINQ method `Where`, finally using `FirstOrDefault` to get the first such entry for further processing. While other languages have equivalent concepts, LINQ is specific to C# and is known to inherent allocations associated with its use. As a result, LINQ usage on the application's hot-path can lead to excessive allocations, which can cause unwanted symptoms such as spikes in garbage

collection (GC), depriving the application of CPU resources, and reducing its throughput.

Figure 4 shows the steps in our approach to generate fixes for a given buggy method and expensive line of code. To fix this code, we first compare the line of code containing the bug (that can be found easily using a profiler [Gayhardt(2019)]), by abstracting out project-specific details and querying a pre-built knowledge-base containing past buggy API usages, which includes the buggy usages of `FirstOrDefault`. This gives us a prompt instruction in English language telling how this bug should be fixed, which we then use to construct the prompt as shown in Figure 2. The prompt includes the buggy method itself as a comment, followed by the prompt instruction, telling the model to re-write the preceding method without the `FirstOrDefault` call, implying that it should be replaced with cheaper alternatives, which in this case would be an inlined `foreach` loop. To further ensure that the the model completion contains the commented method, the prompt also includes the signature of the buggy method, followed by an open curly brace, at the end of prompt. This is to coax the model into outputting a completion of that same method. Figure 3 shows the completion generated by the LLM to this prompt, which indeed replaces the LINQ query with an explicit `foreach` loop.

```
/*
public override void Undo(Params param) {
    foreach (Container obj in containers) {
        Beatmap copy =
            Beatmap.GenerateCopy(obj.objectData);
        param.collections.Where(x => x.Type ==
            copy.Type)
            .FirstOrDefault()?.SpawnObject(copy,
                out _);
        if (obj is EventContainer e &&
            e.eventData.IsRotationEvent)
            param.tracksManager.RefreshTracks();
    }
}
*/
/* PERF: Rewrite the above method without
   FirstOrDefault. */
public override void Undo(Params param) {
```

Fig. 2: Model input prompt used to generate a fix for the method in Figure 1. The prompt consists of (i) the commented original buggy method, (ii) an instruction telling the model how to fix the issue, and (iii) the starting fragment of the buggy method.

III. OUR APPROACH

In this section, we explain the approach behind RAPGen and how we use it at inference time to fix performance bugs. The basis of our approach is building a code transformation knowledge-base with a code change specific retrieval algorithm, which we use to access the relevant prompt instruction for a given buggy code and construct the prompt.

```

public override void Undo(Params param) {
    foreach (Container obj in containers)
    {
        Beatmap copy =
            Beatmap.GenerateCopy(obj.objectData);
        foreach (var collection in param.collections)
        {
            if (collection.Type == copy.Type)
            {
                collection.SpawnObject(copy, out _);
                if (obj is EventContainer e &&
                    e.eventData.IsRotationEvent)
                    param.tracksManager.RefreshTracks();

                break;
            }
        }
    }
}

```

Fig. 3: Suggestion generated by the LLM when asked to complete the prompt in 2. The suggested fix is to unroll the LINQ query in favour of a foreach loop, which can stop early when a matching entry is found. This fix closely matches the developer fix and saves other potential LINQ overheads such as the allocations and GC.




Fig. 4: The fix generation pipeline followed by RAPGen, showing how it is used at inference time to fix a performance issue given a buggy method and corresponding expensive line.

A. Code Transformation Knowledge-base

The first step of our proposed approach is to build a knowledge-base of performance code transformations and corresponding instructions, which can be used to coerce the language model into applying the same code transformation towards fixing a given bug. Below we describe how we collect the relevant data for building such a knowledge-base, as well as the processing steps needed to extract the necessary information to build this knowledge-base.

Data Collection We first collect all repositories with ≥ 5 stars on GitHub, whose primary language is C# and had a commit within the last 5 years. We crawl the commit history of the main branch in each of these repos. A commit typically contains a commit message and a changeset or diff representing the difference between current and previous version of files modified within that commit. From these repos, we collect the set of commits whose titles begin with "PERF:" or "[PERF]". This is a convention [com(2019)] that's often used by developers to indicate that their commit contains primarily performance-related changes. This results in ~ 1500 commits with clean code changes pertaining to performance.

Extracting Performance Bug Patterns. For each commit,

Algorithm 1 Abstract Impact Line

Input: AST *node*, Array *ret*

```

if IsLeaf(node) then
    if IsKeyword(node) or IsSyntax(node) then
        ret = Append(ret, node)
    else if IsCommonIdentifier(node) then
        ret = Append(ret, node)
    else
        ret = Append(ret, Node∅)
    end if
end if
for child ∈ node.Children do
    childArr = AbstractImpactLine(child)
    if ContainsCommonIdentifier(childArr) then
        ret = Concatenate(ret, childArr)
    else
        ret = Append(ret, Node∅)
    end if
end for

```

we take the changeset, $D = [(f_i, f'_i)]_{i=1}^N$, where f_i and f'_i are the before and after versions of the i^{th} file modified in the commit. We parse the code files, f_i and f'_i , and extract all pairs of methods, $[(m_i, m'_i)]_{i=1}^M$, such that m_i and m'_i are the before and after versions of a given method as modified by the commit, respectively.

Given a method from the diff, m_i with K statements $[s_j]_{j=1}^K$, and its corresponding modified version, m'_i with K' statements $[s'_k]_{k=1}^{K'}$, we extract the statements s_j that appear in m_i , but not in m'_i . In other words, we find the set of statements, $S_d = \{s_j \mid 1 \leq j \leq K \text{ and } s_j \notin S_{m'_i}\}$, where $S_{m'_i}$ is the set of statements in m'_i . We then use Algorithm 1 to process all the statements in S_d . The goal of this algorithm is to remove any project-specific identifiers, such as variable names or user-defined function names, replacing them with placeholders. This yields a set of L abstracted performance bug patterns $[p_i]_{i=1}^L$, where $p_i = f(s_j)$ and $s_j \in S_d$.

This leaves us templates of modified usages of common C# API, containing only tokens corresponding to language syntax/keywords and common literals/identifiers in C#. This type of statement abstraction allows us to pick up patterns in changes that get repeated across projects and may correspond to common performance issues. To further avoid any project-specific change patterns, we only keep patterns occur in at least 2 projects.

Below is an example of one such transformation made by the abstraction algorithm:

```

Foo().Where(x => x.Bar()).FirstOrDefault();
      ↓
<Node∅>.Where(<Node∅>).FirstOrDefault();

```

The first line contains a possible statement modified as part of a commit. The abstraction algorithm traverses the AST corresponding to this statement and replaces all project-specific identifiers/literals (e.g. `Foo` and `Bar`) or any subtrees, none of whose descendants are known common identifiers in C#,

with a placeholder node ($\langle Node_{\emptyset} \rangle$). We then traverse the code corresponding to this abstracted tree to get the code with project-specific details abstracted out. In the above example, the identifiers `Where` and `FirstOrDefault` are retained because these are frequently used .NET library functions in C# projects. The above pattern specifically, is an example (Figure 1) of a known inefficient usage pattern where functions, `Where` and `FirstOrDefault`, are chained together inefficiently. This could be re-written to be more optimal by condensing the query to remove the `Where` call or simply unrolling the query into an explicit loop.

Creating a Code Transformation Instruction Now for each extracted performance bug pattern p_i , we find all the K before-after pairs of methods, $[(m_k, m'_k)]$ from performance improvement commits, C_{PERF} , such that the buggy statement, s , which is the concrete version of p_i , is contained within m_k and not in m'_k .

For each such before-after pair, we would like to create an instruction t_k that explains the code transformation taking place as part of the change. To do this, we find $(I_{m_k}, I_{m'_k})$, which are the sets of common C# identifiers added and/or removed as part of this code transformation, respectively.

To extract the identifiers, we compare the sets of statements in m and m' , i.e. S_m and $S_{m'}$, respectively. For every statement in the diff i.e. statements that appear in m and not in m' , or ones that appear in m' and not in m , in other words, $(S_m \setminus S_{m'}) \cup (S_{m'} \setminus S_m)$, we extract any common C# identifiers used within these lines. These common literals/identifiers that correspond to popular library functions in C# and can be found by looking for literals and identifiers that repeat across multiple projects. We consider the set of common C# identifiers found in $(S_m \setminus S_{m'})$ and $(S_{m'} \setminus S_m)$. We then remove any identifiers that are found to be common among the two sets. We call the resulting disjoint sets of identifiers I_m and $I_{m'}$, respectively.

To create t_k , we focus on the following code fix transformations that involve adding or removing identifiers (or APIs):

- **Identifier Addition:** In this case, $I_{m_k} = \emptyset$ and $I_{m'_k} \neq \emptyset$ and $t_k = \text{"PERF: Rewrite the above method with } \langle X \rangle \text{"}$. This represents an instruction to re-write the code with some additional identifiers (or APIs) needed to resolve the performance issue, which, in terms of code-changes, may be a new function call or an additional parameter added to an existing call, etc.
- **Identifier Removal:** In this case, $I_{m_k} \neq \emptyset$ and $I_{m'_k} = \emptyset$ and $t_k = \text{"PERF: Rewrite the above method without } \langle X \rangle \text{"}$. This represents an instruction to remove the usage of some identifiers (or APIs), which may correspond to removal of the use of some inefficient APIs, within the code. The example in Figure 2 falls within this category.
- **Identifier Addition and Removal:** In this case, $I_{m_k} \neq \emptyset$ and $I_{m'_k} \neq \emptyset$ and $t_k = \text{"PERF: Use } \langle X \rangle \text{ instead of } \langle Y \rangle \text{ in the above method."}$. This instruction is meant to represent the case where the model should replace one set of identifiers (or APIs) with another, i.e., situations where one set of API must be swapped out for another, which are more appropriate for the situation.

In each case, we replace the placeholders in the template t_k , i.e. $\langle X \rangle$ and $\langle Y \rangle$, with a string containing the comma-delimited list of identifiers in I_m and $I_{m'}$, respectively.

Knowledge-base Creation Finally, taking all the buggy usage patterns and their corresponding before-after pairs and instructions, we create a knowledge-base to store all the $(p, (m, m', t))$, where p serves as the lookup key to find the set of triples, (m, m', t) , where m and m' are the before-after methods, respectively, and t is the derived instruction from this particular before-after pair. This knowledge-base is essentially a mapping between abstracted performance bug patterns and instructions on appropriate code transformations to fix those bugs. In the following section, we describe our retrieval process, to fetch the right instruction from the knowledge-base.

B. Generating the Prompt

We are given buggy method, m_b , containing an expensive line of code, l , which causes the performance bug. We first abstract the expensive line using Algorithm 1 and get the abstracted pattern p_b . We then perform a lookup in our knowledge-base using the p_b , as the query. This results in a list of K_p code change and instruction triples, i.e. $[(m_k, m'_k, t_k)]_{k=1}^{K_p}$, where m_k and m'_k are the before and after versions of a method, where a similar buggy line was fixed and t_k is the code transformation instruction constructed based on changes made to the identifiers in the before-after pair. Since we want to find the entry that's closest to m_b , we use a code-search technique [Luan et al.(2019)] to rank each before method, m_i , based on similarity to m_b . We take the entry corresponding to the most similar before method, (m_j, m'_j, t_j) and construct the prompt using the retrieved code transformation instruction t_j .

Now we fill the template in Figure 5 to build the final prompt, which will be given to the model. We include the buggy method, m_b , itself surrounded by C-style comments, followed by t_j and the signature of m_b and an open curly-brace, which in C# indicates the start of the method body.

```
/*
<Commented Buggy Method>
*/
/* PERF: <Prompt Message> */
<Method Signature> {
```

Fig. 5: High level prompt template followed by the prompts generated by RAPGen. They consist of the buggy method itself, followed by the prompt instruction retrieved from the KB using the buggy line within the buggy method and finally the signature of the method itself, proceeded by an open curly brace.

Figure 2 shows a concrete example of one such prompt following the above template.

C. Generating Fixes

Generating the fix itself is relatively simple. We use the constructed prompts to generate completions using the *gpt-3.5*.

turbo [ope(2023)] model through OpenAI’s REST API. Since it is difficult to tell the model when to stop, it will likely go on outputting code even after it has generated the fixed method. We parse out the completion corresponding to the method by extracting the code until curly braces are balanced and discard the rest of the output.

IV. EMPIRICAL EVALUATION

In this section, we present our empirical evaluation we conducted on understanding our prompt design choices and assessing the effectiveness of our approach. We first describe our experimental setup, i.e. our dataset, metrics and baselines for our empirical evaluation.

A. Experimental Setup

Dataset We leverage the dataset containing past performance fixes from the DeepDev-PERF study [Garg et al.(2022)]. This dataset consists of 132 instances of performance fixes, which were manually reviewed and confirmed to be performance-related changes by performance experts after examining a broad set of C# commits on GitHub. Furthermore, it was shown that this dataset covers a wide-range of performance bugs. Therefore, a good accuracy over this dataset would demonstrate the effectiveness of a model for this task. Each example in this dataset is of the form (m, m') , where m and m' are the before and after versions, respectively. These are also mostly single method performance improvement fixes, making them ideal for our scenario.

Adding Bug Localization Information Our approach assumes that bug localization information, i.e. the expensive line, is included as part of the input. In practice, one can leverage performance bug localization using profiling data to find the expensive line [Garg et al.(2021)]. However, for the purpose of evaluating our approach, we used a heuristic to localize bugs in the above dataset. We examine the diff between the before and after methods, m and m' and take the first line $l_b \in m$ that does not match the corresponding line in m' , when compared in order, as the buggy line.

Inference Hyperparameters We generate 100 suggestions/fixes for each example in the dataset, using a relatively large temperature of 0.7 to achieve sufficient variety within the fixes generated. We also use a max token limit of 1024 tokens, so that the model has enough token bandwidth to generate the entire method for longer suggestions. Finally, we take the model output and parse out the fix as described in the previous section.

Automated Metrics We use the following three automated metrics for measuring the effectiveness of our approach:

- **Verbatim Match %:** This metric considers a prediction to be correct only if the generated fix exactly matches the human fix in the test set. This metric usually presents a lower bound as two pieces of code can be equivalent, despite not matching verbatim. This metric is also sensitive to identifier names and order of code statements.
- **Abstracted Match %:** In this metric, we abstract out variable names using a placeholder name (such as

"VAR_{i}", e.g. "VAR_0", "VAR_1", "VAR_2" ...), where i is determined based on the order, in which we encounter the variables during a traversal of the AST of the method. Using the traversal order, ensures that two code snippets with the same control flow and no major changes except variable names, will have the same abstracted variable names. But, even this presents us with a lower bound on the approach’s true capabilities in fixing bugs as a fix may have different control-flow from ground truth, but still be correct.

- **CodeBLEU:** Our final metric is CodeBLEU [Ren et al.(2020)], which is a variation of the popular NLP metric, BLEU [Papineni et al.(2002)]. In addition to n-gram matching, it also takes code characteristics like abstract syntax tree (AST) and data-flow similarity into account, when comparing programs. Similar to the DeepDev-PERF study, we use the hyper-parameters that had the highest correlation to human scores in the CodeBLEU study, i.e. $\alpha, \beta, \gamma, \delta = 0.1, 0.1, 0.4, 0.4$.

Human Evaluation Metrics In addition to the above automated metrics, we also conduct a human evaluation of the suggestions with two performance experts (not on the author list). We use the following metric:

- **Closest Match Top-K Accuracy %** Due to there being too many suggestions to go through one by one, we use the state of the art code search technique, Aroma [Luan et al.(2019)], to find the suggestion closest to the ground truth. The most similar suggestion is shown to two performance experts, who are then asked whether they consider the model suggestion to be equivalent or better than the developer fix in terms of performance. We consider a suggestion to be correct, if it is found to be correct by both experts. This allows us to get a lower-bound on the model’s actual Top-K accuracy, if one were to exhaustively go through every suggestion.

Prompt Variants To better understand our prompt design choices, we investigate the following prompt variations in addition to RAPGen:

- **Static Prompt Message.** This prompt follows the template in Figure 6a. We first add the buggy method, m_b itself as a C-style comment, followed by a static comment expressing an intent to improve performance and, finally, the signature of the buggy method with an open curly brace. The purpose behind including this prompt is to contrast with our approach where we provide a prompt message containing information regarding where the bug is and hints on how to fix it, to show whether providing the model with additional information regarding the bug helps it come up with more fixes.
- **Retrieval-based One-shot Prompt.** This prompt follows the structure in Figure 6b. We first provide the model with the retrieved before-after pair, (m_j, m'_j) (from Section III-B) where m_j contains a line, l_{m_j} , whose abstracted version matches the abstracted buggy line, l_b in m_b . Similar to

```

/*
<Commented Buggy Method>
*/
/* PERF: Improve performance of
the above method. */
<Buggy Method Signature> {

```

(a) Static Prompt Template

```

/*
<Commented Retrieved Buggy Method>
*/
<Retrieved Fixed Method>
/*
<Commented Buggy Method>
*/
<Buggy Method Signature> {

```

(b) Retrieval-based One-Shot Prompt Template

```

/*
<Commented Buggy Method>
*/
/* PERF: <X> is on the
hot-path in the above
method. We can do better
by ...

```

(c) Reasoning-based Prompt Template

Fig. 6: Prompt templates for various prompt variations we tried in our evaluation.

RAPGen prompts, we then provide the model with the buggy method as a comment, followed by the signature of the buggy method, coaxing it into outputting a fixed version of that method. Our reason for including this prompt is to show the impact of providing a raw fix example, as opposed to giving the model the instruction derived from the same code-change.

- *Reasoning-based Prompt Message.* Recent work on prompt engineering suggests letting LLMs generate an intermediate reasoning about a problem can greatly improve performance on a given task [Wei et al.(2022)]. As our final variant, we designed a reasoning-based prompt that follows the structure in Figure 6c. This prompt also includes the buggy method, m_b , followed by an incomplete comment telling the model what call is being expensive, and coaxing the model to reason about the problem and then generate a possible fix. The main idea is that we first get the model to reason about the problem and come up with a description of the possible fix by completing the prompt message itself, before it generates the fixed code. We do this by giving the model the buggy method, followed by the prompt message. We let the model complete the rest of the comment (i.e. continuing in place of the ellipsis) and then have it generate a fix as well. The prompt follows structure in Figure 6c.

B. Comparing with Other Prompt Variants

Table I shows the results of our automated and human evaluation metrics, achieved by RAPGen and other prompt variants. Our approach is able to achieve a $\sim 42\%$ verbatim match and a $\sim 41\%$ abstracted verbatim match score over the dataset, which is higher than any of the other prompt variants. We also see that in the human evaluation, our approach is able to suggest valid improvements for up to $\sim 63\%$ suggestions within just 100 sampling attempts, higher than all our other

prompt variant. We discuss these results and possible reasons behind them in more detail below.




Fig. 7: Closest Match Top-K Accuracy plot of our prompts on the test set. We can see that RAPGen achieves the best score among the 4 prompt variants.

Effect of adding instruction conveying the bug location and possible fix. We first compare the results of the RAPGen’s prompts with the Static prompts (Figure 6a). These two prompts are almost identical except for the prompt message following the commented buggy method. RAPGen’s prompt includes a fix instruction based on past fixes to bugs of similar nature, which tells the model what function call within the method body is expensive, as well as a possible fix. On the other hand, the Static Prompt uses the same generic fix instruction for each bug. We see that RAPGen performs significantly better in terms of both the number of exact matches it finds, $>10\%$ more, as well as, achieves a higher Top-K score (by $>25\%$) compared to static prompts. Figure 9 shows an example of

TABLE I: Summary of the results of our approach against other prompt variants.

Model	Verbatim Match %	Abstracted Match %	CodeBLEU	Closest Match Top-K Accuracy %		
				1	10	100
Static Prompt	26.5	27.3	67.6	4.5	13.6	33.3
Retrieval-based One-shot Prompt	28.8	29.5	69.7	6.1	14.4	37.9
Reasoning-based Prompt	21.2	31.1	53.6	1.5	9.1	43.2
RAPGen	42.4	45.5	72.0	17.4	28.8	60.6

TABLE II: Summary of the results of our approach against other models.

Model	Verbatim Match %	Abstracted Match %	CodeBLEU	Closest Match Top-K Accuracy %		
				1	10	100
DeepDev-PERF	34.3	37.3	70.7	8.3	18.2	34.1
RAPGen	42.4	45.5	72.0	17.4	28.8	60.6

a case where RAPGen was able to find a fix, whereas the Static prompt wasn't. Since this was a relatively long method, which required a non-trivial fix consisting of a data-structure change, Static prompt fails to come up with a fix. Whereas, RAPGen prompt is guided by the prompt instruction based on past fixes, which is better able to guide the model to the right fix. Intuitively, this makes sense because RAPGen's prompt message includes information about the location of the bug as well.

Effect of extracting instruction from retrieved example.

Prior work has shown that providing an example in a few shot setting usually results in better performance in a downstream task than providing an instruction [Brown et al.(2020)]. It likely has less to do with the example but more with specificity of the instructions. The Retrieval-based One-Shot prompt gives the model the same retrieved before-after pair that the instruction is derived from. However, comparing the results of RAPGen with the One-Shot prompt, we can see that going the extra step of extracting an instruction from the retrieved before-after pair improves the results significantly. We believe that this may be because, given the raw before-after method pair, the model has to perform the diff implicitly and infer the instruction by itself, before applying the change to the unseen method. It is possible for it misunderstand the change and make mistakes in understanding the nature of the fix. We see that our prompts find more exact matches (by $>10\%$) and achieves a higher Top-K accuracy score (by $>20\%$), compared to retrieval prompts. For the example shown in Figure 9, both the One Shot and RAPGen prompts were able to generate the right fix in our evaluation. However, the One Shot prompt took almost twice as many attempts before arriving at the right fix, compared to RAPGen. The likely reason RAPGen is able to generate a correct suggestion more quickly may be due to a possible fix being readily available as part of the instruction, without the model having to perform in-context learning from a raw code change example and infer the underlying transform.

Effect of letting the model reason about the problem on its own. Both Reasoning-based and RAPGen's prompts include information about where the buggy line is within the method. The main difference between them is that we do not provide the model with any hints on how to fix the bug in Reasoning-based prompts and instead allow it to reason about the problem and come up with a possible fix itself. During the

human evaluation, we noticed that with such a reasoning-based prompt, the model is more likely to come up with high-level changes (introducing caching, memoization, adding a fast-path, etc.), modifying the control flow or the algorithm itself rather than an API-level fix. Figure 8 shows an example of one such change suggested for a reasoning based prompt, during our evaluation. Since high-level changes can be more complex to make, the model is more likely to create an incorrect or over-complicated fix. Therefore, it cannot arrive at the correct fix right away resulting in a flatter plot in the Figure 7, despite given the problematic function call.

C. Comparison with State-of-the-art

Finally, we compare our approach with the current state of the art approach for performance bug fixing for C#: DeepDev-PERF [Garg et al.(2022)]. DeepDev-PERF uses a sequence to sequence transformer model, BART [Lewis et al.(2020)] and finetunes it for the task of generating performance bugs. The model can output entire patches, as opposed to just completions, allowing it to generate fixes spanning the entire class. In contrast, our approach is intended to fix single method bugs only. Therefore, we only compare our model with DeepDev-PERF's ability to fix single method fixes, using the dataset they provide.

Table II shows that RAPGen outperforms DeepDev-PERF both in automated and human evaluation metrics. This result is quite encouraging given that RAPGen only leverages our novel prompting approach in zero-shot, while DeepDev-PERF relies on expensive fine-tuning steps over a large dataset. Furthermore, it is also important to note that in DeepDev-PERF, the authors allow their model 2000 attempts for each example, and use the top-500 (based on average token likelihoods) for their evaluation, whereas, we only sample 100 completions per example. Despite the much fewer sampling attempts, our approach is still able to out-perform their finetuned model. One reason for this gap may be that our model includes information regarding buggy line, within the prompt, whereas the input to DeepDev-PERF doesn't localize the bug beyond just the buggy method. Therefore, their model may need to suggest many possible fixes before it finds the right one.

V. IN-THE-WILD EVALUATION

To see the effectiveness of our approach in practice, we conducted an evaluation on existing .NET projects to verify

```

/*
public IServiceProvider
CreateContainer(ShellSettings settings)
{
    // ...
    foreach (var dependency in
        blueprint.Dependencies)
    {
        foreach (var interfaceType in
            dependency.GetInterfaces())
        {
            _logger.LogDebug("Type: {0},
                Interface Type: {1}",
                dependency.Type, interfaceType);
        }
    }
    // ...
}
*/
/* PERF: LogDebug is on the hot-path in the
above method. We can do better by ...

```

```

using LogDebug only when logging is enabled. */
public IServiceProvider
CreateContainer(ShellSettings settings)
{
    foreach (var dependency in
        blueprint.Dependencies)
    {
        foreach (var interfaceType in
            dependency.GetInterfaces())
        {
            if(!_logger.IsEnabled(LogLevel.Debug))
            {
                _logger.LogDebug("Type: {0},
                    Interface Type: {1}",
                    dependency.Type, interfaceType);
            }
        }
    }
    // ...
}

```

Fig. 8: A relatively high-level suggestion generated by the Reasoning-based prompt. The code on top (simplified) contains the commented out buggy method, followed by an incomplete comment letting the model know about where the problem is, which, in this case, happens to be excessive calls to `LogDebug`. The suggestion at the bottom finishes the incomplete the comment with a message, showing that the model did indeed reason about the problem presented and came up with a possible fix. The suggestion that follows wraps the `LogDebug` calls with a `Logger.IsEnabled` checks. This was also the fix made by the developer.

the effectiveness of our approach in helping developers fix real world performance problems.

A. Experimental Setup

1) *Finding real world problems to fix:* We start with 50 codebases, internal to a large company, that are associated with existing services running in the cloud. Each of these services is being profiled by Microsoft Application Insights Profiler [Gayhardt(2019)]. The collected profiles are then analyzed by a performance bug detection service called Code Optimizations [cod(2023)] that continuously analyzes profiler traces from

```

/*
private RectangleF CreateBBox()
{
    List<Vertex> processed = new List<Vertex>();
    foreach (Sidedef s in sidedefs)
    {
        if (!processed.Contains(s.Line.Start))
        {
            if (s.Line.Start.Position.x < left)
                left = s.Line.Start.Position.x;
            // ...
            processed.Add(s.Line.Start);
        }
    }
    // ...
}
*/
/* PERF: Use ContainsKey instead of Contains in the above method. */
private RectangleF CreateBBox()
{

```

```

    Dictionary<Vertex, bool> processed = new Dictionary<Vertex, bool>();
    foreach (Sidedef s in sidedefs)
    {
        if (!processed.ContainsKey(s.Line.Start))
        {
            if (s.Line.Start.Position.x < left)
                left = s.Line.Start.Position.x;
            // ...
            processed[s.Line.Start] = true;
        }
    }
    // ...
}

```

Fig. 9: An API change suggested by RAPGen for one of the examples in the testset. The code on the top (shortened for brevity) shows the commented out buggy method with performance issues caused by the use of `List.Contains()` (a $O(N)$ search) within a for-loop. The instruction retrieved from the knowledge-base is to replace `List.Contains()` with a `Dictionary.ContainsKey()` (an $O(1)$ lookup). RAPGen prompt also includes the signature of the method, followed by an open curly brace to coax the model into outputting the same method. The second code snippet shows the completion generated by the model for this prompt, which is indeed to replace the `List` instantiation with that of a `Dictionary` and updating the calls within the loop to use `ContainsKey`. This change was also semantically equivalent to the fix made by the developer.

a given application in the cloud and finds performance bottlenecks.

Each bottleneck is of the form (b, p) , where b in the bottleneck method that’s usually an API method and p the calling method from the user’s codebase. We look at the insights generated by Code Optimizations over CPU profiler traces. For each service, we find bottlenecks (b, p) that meet the following criteria:

- The bottleneck takes up more than $>5\%$ of the application’s CPU. We set this threshold because we want to fix issues that take up a substantial amount of resources and are worth fixing. These are the issues that a developer would typically look into when improving the performance of their application.
- b is present in the RAPGen knowledgebase (Section III-A)
- A call to the method b is present in the code for the parent method p . For example, if the bottleneck method b is `Dictionary.ContainsKey`, we look for a call to `ContainsKey` in the parent method body. We need

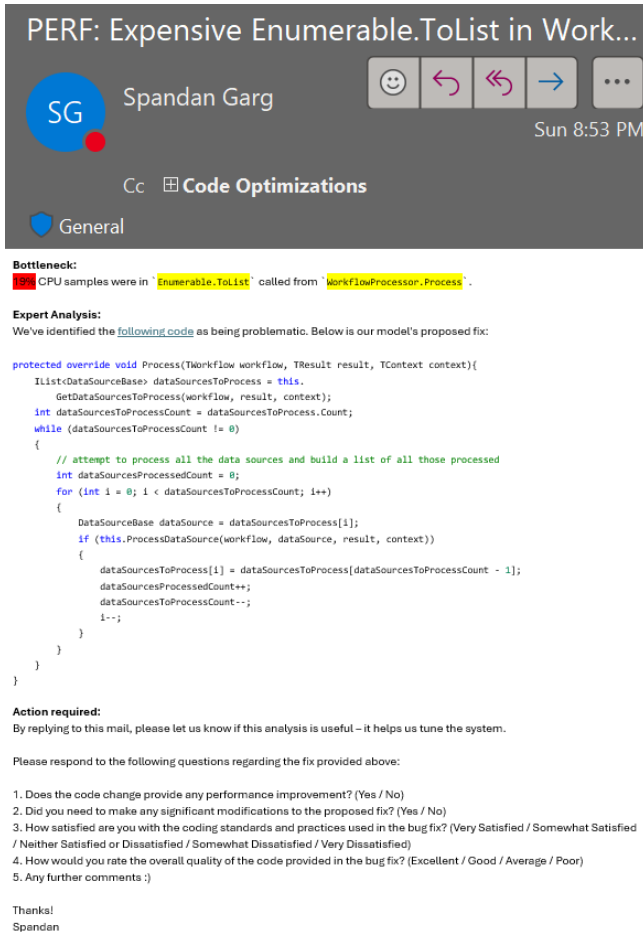


Fig. 10: An email sent as part of our "In-the-wild" Evaluation. The email contains a fix for a performance bottleneck found using Code Optimizations [cod(2023)] in a Microsoft service. Along with the fix, we also include information regarding the bottleneck, i.e. the bottleneck method and the calling parent method (both highlighted in yellow) and current CPU usage of the bottleneck. At the end of the email, we ask the users to answer a few questions regarding the fix (Section V-A2).

to confirm this because, this is often not the case when there is inlining of methods or the code uses syntactic sugar (e.g. use of special characters like "+" to perform `String.Concat` when concatenating strings). III-A)

2) *Reaching out to developers*: We take 10 random bottlenecks that meet the above criteria and each come from a different service. We then locate the code associated with each bottleneck in the corresponding codebase and use RAPGen to generate a fix. We communicate the generated fix to the engineers working on the project via email (Figure 10). In the email, we ask them to respond to the following questions regarding the fix:

- (Q1) Does the code change provide any performance improvement? (Yes / No)
- (Q2) Did you need to make any significant modifications to the proposed fix? (Yes / No)

- (Q3) How satisfied are you with the coding standards and practices used in the bug fix? (Very Satisfied / Somewhat Satisfied / Neither Satisfied or Dissatisfied / Somewhat Dissatisfied / Very Dissatisfied)
- (Q4) How would you rate the overall quality of the code in the bug fix? (Excellent / Good / Average / Poor)

B. Results

We were able to secure responses to 8 of the emails i.e. a 80% response rate. 7 of the 8 respondents confirmed that the changes would improve the performance of the code. 1 of the respondents rejected the fix. 6 in 7 respondents who accepted the fix confirmed that the changes could be accepted without any modifications. One of the respondents who accepted needed to make a minor change to the code.

For the next question, we asked the developer to rate their satisfaction by picking from the following options: Very Satisfied, Somewhat Satisfied, Neither Satisfied or Dissatisfied, Somewhat Dissatisfied and Very Dissatisfied. 6 in 7 developers responded with *Very Satisfied*. The final developer responded with *Satisfied* because they had to make slight modifications to the code. As our final question, we asked the developer to pick from the following options: Excellent, Good, Average, Poor. 6 in 7 developers rated the overall quality of the fix as *Excellent*, except for one developer who responded with *Good*. Overall, this in-the-wild evaluation confirms the usefulness of our approach for fixing perf bugs in real life applications.

VI. THREATS TO VALIDITY

In this work, we focus only on the subset of performance issues caused by misuse of common .NET APIs. However, it's possible to have performance problems that are due to non-API related reasons [Garg et al.(2022)], such as unnecessary computation in a loop that could be hoisted out or excessive allocations of a given type, etc. We also limit our scope to performance issues that can be fixed by changes to a single file and class. However, it's possible for fixes to span multiple classes, files or folders. This would require a much larger context window as we would need to pass code from multiple classes and the model would then decide which need to be modified in order to fix the issue. A concrete example of this could be a fix that requires making changes to a class that's used in multiple locations in a given code-base. The fix may require changes to not only the problematic class, but the usages as well. We plan to explore this in future work.

Additionally, we only explore fixing performance issues and assume that there already exists a way to identify expensive parts of a code-base, down to the exact line number. In .NET, this is indeed possible. One could use a profiler [Gayhardt(2019)] and .NET's Source Link [sou(2018)] technology, which adds metadata containing line number, commit, branch, etc. to a given profiler trace. We have not yet explored this for other languages like Java, python, etc., but we assume that similar technologies exist there as well. We leave this exploration to future work.

Finally, we were unable to run the code fixes suggested to the benchmark dataset. This is because many of the GitHub projects these examples were drawn from were not buildable, for variety of reasons such as, requiring proprietary packages to build, custom build steps or simply compiler errors. However, since our human evaluation showed that the fixes were equivalent to those proposed by the developer, we believe that our changes would result in performance gain as well. In practice, one could add an additional step to run benchmark tests to verify the change improves performance. Furthermore, the benchmark test could itself be generated using an LLM. We leave this exploration to future work. There already exists work that has already explored generating unit tests using these models [Tufano et al.(2020)].

VII. RELATED WORK

Next we discuss the work closely related to ours.

A. Automated Program Repair

Traditional automated program repair (APR) techniques usually require test suits or logic assertions as a specification of correctness. For example, GenProg [Weimer et al.(2009)] uses genetic programming to generate code fixes that pass a supplementary suite of test cases. Pattern-based Automatic program Repair (PAR) [Kim et al.(2013)] leverages manually created fix templates based on patterns in existing developer patches to guide the genetic search. Similarly, CapGen [Wen et al.(2018)] attempts to generate patches at finer granularity (AST node level) using context-aware prioritization of genetic mutation operations, and VarFix [Wong et al.(2021)] further extends GenProg by using variational execution [Austin and Flanagan(2012)]. One of the major challenges with these APR approaches is overfitting [Ye et al.(2021)] and high maintenance cost. They also usually require a significant amount of effort on modeling program semantics. To fill this gap, our approach leverages prompt engineering on LLMs and therefore fixes a wide-range of performance bugs with low training and maintenance cost.

B. Neural Code Fix Generation

Recent advancements in deep learning, which automatically learn patterns from unstructured data, have driven a shift in APR techniques away from genetic programming and towards learning-based methods. Machine learning-based APR approaches leverage generative deep learning models that directly output patched versions of faulty code. For example, [Tufano et al.(2018)] uses an RNN encoder-decoder model to generate fixes for bugs on GitHub projects. DeepDebug [Drain et al.(2021)] trained a backtranslation model to generate synthetic buggy versions of code to fine-tune a bug patch model. Stack traces for these generated bugs, as well as supplementary context through code skeletons, are used to finetune a bug patch generation model. Similarly, VRepair [Chen et al.(2022)] trains a vanilla Transformer to automatically repair security vulnerabilities using transfer learning. The model is initially pretrained on bug-fix pairs from GitHub and fine-tuned on

security vulnerability-fix pairs. VulRepair [Fu et al.(2022)] builds on VRepair by replacing the vanilla Transformer with a T5 [Raffel et al.(2020)] architecture with BPE tokenization. TFix [Berabi et al.(2021)] is a transformer-based model that fixes bugs found with a static analyzer. More recently, studies have shown that Codex [Chen et al.(2021)] can also generate patches to vulnerable code snippets [Pearce et al.(2021a)]. Similarly, [Prenner et al.(2022)] explored the viability of using Codex in a zero-shot setting to fix bugs in the QuixBugs benchmark [Lin et al.(2017)]. In this setting, Codex was indeed able to generate bug fixes, although fix performance was heavily dependent on the prompts fed to the model, which were created manually. While the majority of the work is focused on general bug fixing or vulnerability fixes, our work uses prompt engineering for fixing performance bugs specifically.

C. Prompting and Prompt Engineering

With the rise of LLMs prompt engineering [Liu et al.(2023)] has emerged as a process to come up with the best prompt for a downstream task. A high quality prompt is often necessary to elicit the correct response from a language model. However, coming up with the optimal prompt is no trivial task [Jiang et al.(2020)]. Prior work have explored various prompt engineering approaches including manually engineering suitable prompts [Petroni et al.(2019)], [Brown et al.(2020)], few-shot prompting, and optimizing the ordering of prompt examples [Lu et al.(2021)] among others. Similarly, continuous prompts have emerged as an automated prompt learning approach [Li and Liang(2021)], [Liu et al.(2021)], [Lester et al.(2021)]. Continuous prompts directly prompt the embedding space of the model and usually require labeled data. Our work uniquely contributes to prior work by introducing a new prompt generation method called RAPGen, which leverages a knowledge-base of instructions to guide the generation of the prompt.

VIII. CONCLUSION

In this work, we presented Retrieval Augmented Prompt Generation (RAPGen), a novel prompt engineering approach for fixing performance bugs using LLMs. RAPGen uses past performance fixes to construct a knowledge-base instructions encoding the changes made within these fixes. When given a code snippet with a performance issue and the corresponding expensive line, RAPGen leverages this knowledge-base to construct a prompt and generates fixes using LLMs in zero-shot. Being zero-shot and entirely prompt-driven enables RAPGen to leverage the capabilities of LLMs for Code without the cost of fine-tuning. Our extensive empirical evaluation shows that about 40% of the fixes produced by RAPGen match the developer's fix and in ~60% of the cases, it is able to generate a fix that is equivalent or better than the developer fix. Our "in-the-wild" evaluation shows that the model is capable of producing fixes for performance problems in real world applications that are considered useful by developers. While we only explored our approach for the task of performance bug fixing, our knowledge-base creation and retrieval methods

are not specific to performance bugs and can be extended to other types of bugs as well. We believe that our work uniquely contributes to the area of prompt engineering, specially for code-related downstream tasks.

REFERENCES

- [sou(2018)] 2018. SourceLink. <https://github.com/dotnet/sourcelink>
- [com(2019)] 2019. Conventional Commits. <https://www.conventionalcommits.org/en/v1.0.0/>
- [ope(2023)] 2023. Azure OpenAI Models. <https://learn.microsoft.com/en-us/azure/cognitive-services/openai/concepts/models>
- [cod(2023)] 2023. Code Optimizations. <https://techcommunity.microsoft.com/t5/azure-observability-blog/announcing-the-public-preview-of-code-optimizations/ba-p/3824757>
- [Attariyan et al.(2012)] Mona Attariyan, Michael Chow, and Jason Flinn. 2012. X-ray: Automating {Root-Cause} Diagnosis of Performance Anomalies in Production Software. In *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*. 307–320.
- [Austin and Flanagan(2012)] Thomas H. Austin and Cormac Flanagan. 2012. Multiple Facets for Dynamic Information Flow. *SIGPLAN Not.* 47, 1 (jan 2012), 165–178. <https://doi.org/10.1145/2103621.2103677>
- [Berabi et al.(2021)] Berkay Berabi, Jingxuan He, Veselin Raychev, and Martin Vechev. 2021. Tfix: Learning to fix coding errors with a text-to-text transformer. In *International Conference on Machine Learning*. PMLR, 780–791.
- [Bielik et al.(2017)] Pavol Bielik, Veselin Raychev, and Martin Vechev. 2017. Learning a static analyzer from data. In *International Conference on Computer Aided Verification*. Springer, 233–253. <https://doi.org/10.48550/arXiv.1611.01752>
- [Brown et al.(2020)] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [Chen et al.(2021)] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde, Jared Kaplan, Harrison Edwards, Yura Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, David W. Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William H. Guss, Alex Nichol, Igor Babuschkin, S. Arun Balaji, Shantanu Jain, Andrew Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew M. Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. *ArXiv abs/2107.03374* (2021).
- [Chen et al.(2019)] Yiqun Chen, Stefan Winter, and Neeraj Suri. 2019. Inferring Performance Bug Patterns from Developer Commits. In *2019 IEEE 30th International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 70–81.
- [Chen et al.(2022)] Zimin Chen, Steve James Komrmusch, and Martin Monperrus. 2022. Neural Transfer Learning for Repairing Security Vulnerabilities in C Code. *IEEE Transactions on Software Engineering* (2022), 1–1. <https://doi.org/10.1109/TSE.2022.3147265>
- [Dean et al.(2014a)] Daniel J Dean, Hiep Nguyen, Xiaohui Gu, Hui Zhang, Junghwan Rhee, Nipun Arora, and Geoff Jiang. 2014a. PerfScope: Practical online server performance bug inference in production cloud computing infrastructures. In *Proceedings of the ACM Symposium on Cloud Computing*. 1–13.
- [Dean et al.(2014b)] Daniel J. Dean, Hiep Nguyen, Xiaohui Gu, Hui Zhang, Junghwan Rhee, Nipun Arora, and Geoff Jiang. 2014b. PerfScope: Practical Online Server Performance Bug Inference in Production Cloud Computing Infrastructures. In *Proceedings of the ACM Symposium on Cloud Computing (Seattle, WA, USA) (SOCC '14)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/2670979.2670987>
- [Della Toffola et al.(2015)] Luca Della Toffola, Michael Pradel, and Thomas R. Gross. 2015. Performance Problems You Can Fix: A Dynamic Analysis of Memoization Opportunities. *SIGPLAN Not.* 50, 10 (oct 2015), 607–622. <https://doi.org/10.1145/2858965.2814290>
- [Drain et al.(2021)] Dawn Drain, Colin B. Clement, Guillermo Serrato, and Neel Sundaresan. 2021. DeepDebug: Fixing Python Bugs Using Stack Traces, Backtranslation, and Code Skeletons. *ArXiv abs/2105.09352* (2021). <https://doi.org/10.48550/arXiv.2105.09352>
- [Fu et al.(2022)] Michael Fu, Chakkrit Tantithamthavorn, Trung Le, Trung Le, Van Nguyen, and Dinh Phung. 2022. VulRepair: A T5-Based Automated Software Vulnerability Repair. <https://doi.org/10.1145/3540250.3549098>
- [Garg et al.(2022)] Spandan Garg, Roshanak Zilouchian Moghaddam, Colin B. Clement, Neel Sundaresan, and Chen Wu. 2022. DeepDev-PERF: A Deep Learning-Based Approach for Improving Software Performance (*ESEC/FSE 2022*). Association for Computing Machinery, New York, NY, USA, 948–958. <https://doi.org/10.1145/3540250.3549096>
- [Garg et al.(2021)] Spandan Garg, Roshanak Zilouchian Moghaddam, Neel Sundaresan, and Chen Wu. 2021. PerfLens: a data-driven performance bug detection and fix platform. In *Proceedings of the 10th ACM SIGPLAN International Workshop on the State Of the Art in Program Analysis*. 19–24.
- [Gayhardt(2019)] Lauryn Gayhardt. 2019. *app insights overview*. <https://docs.microsoft.com/en-us/azure/azure-monitor/app/app-insights-overview>
- [Iqbal et al.(2021)] Md Shahriar Iqbal, Rahul Krishna, Mohammad Ali Javidian, Baishakhi Ray, and Pooyan Jamshidi. 2021. CADET: Debugging and Fixing Misconfigurations using Counterfactual Reasoning. <https://doi.org/10.48550/arXiv.2010.06061>
- [Jiang et al.(2020)] Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. How Can We Know What Language Models Know? *Transactions of the Association for Computational Linguistics* 8 (07 2020), 423–438. https://doi.org/10.1162/tac1_a_00324 *arXiv:https://direct.mit.edu/tac1/article-pdf/doi/10.1162/tac1_a_00324/1923867/tac1_a_00324.pdf*
- [Jovic et al.(2011)] Milan Jovic, Andrea Adamoli, and Matthias Hauswirth. 2011. Catch Me If You Can: Performance Bug Detection in the Wild. *SIGPLAN Not.* 46, 10 (oct 2011), 155–170. <https://doi.org/10.1145/2076021.2048081>
- [Kim et al.(2013)] Dongsun Kim, Jaechang Nam, Jaewoo Song, and Sunghun Kim. 2013. Automatic patch generation learned from human-written patches. In *2013 35th International Conference on Software Engineering (ICSE)*. 802–811. <https://doi.org/10.1109/ICSE.2013.6606626>
- [Lester et al.(2021)] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691* (2021).
- [Lewis et al.(2020)] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 7871–7880. <https://doi.org/10.18653/v1/2020.acl-main.703>
- [Li and Liang(2021)] Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190* (2021).
- [Lin et al.(2017)] Derrick Lin, James Koppel, Angela Chen, and Armando Solar-Lezama. 2017. QuixBugs: A multi-lingual program repair benchmark set based on the Quixey Challenge. In *Proceedings Companion of the 2017 ACM SIGPLAN international conference on systems, programming, languages, and applications: software for humanity*. 55–56.
- [Liu et al.(2023)] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *Comput. Surveys* 55, 9 (2023), 1–35.
- [Liu et al.(2021)] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021. GPT understands, too. *arXiv preprint arXiv:2103.10385* (2021).
- [Lu et al.(2021)] Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786* (2021).
- [Luan et al.(2019)] Sifei Luan, Di Yang, Celeste Barnaby, Koushik Sen, and Satish Chandra. 2019. Aroma: Code recommendation via structural code search. *Proceedings of the ACM on Programming Languages* 3, OOPSLA (2019), 1–28. <https://doi.org/10.1145/3360578>
- [Nistor et al.(2013)] Adrian Nistor, Tian Jiang, and Lin Tan. 2013. Discovering, reporting, and fixing performance bugs. *2013 10th Working*

- Conference on Mining Software Repositories (MSR)* (2013), 237–246. <https://doi.org/10.1109/MSR.2013.6624035>
- [Olivo et al.(2015)] Oswaldo Olivo, Isil Dillig, and Calvin Lin. 2015. Static detection of asymptotic performance bugs in collection traversals. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 369–378.
- [Papineni et al.(2002)] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Annual Meeting of the Association for Computational Linguistics*.
- [Pearce et al.(2021a)] Hammond Pearce, Benjamin Tan, Baleegh Ahmad, Ramesh Karri, and Brendan Dolan-Gavitt. 2021a. Examining Zero-Shot Vulnerability Repair with Large Language Models. <https://doi.org/10.48550/ARXIV.2112.02125>
- [Pearce et al.(2021b)] Hammond A. Pearce, B. Tan, Baleegh Ahmad, Ramesh Karri, and Brendan Dolan-Gavitt. 2021b. Can OpenAI Codex and Other Large Language Models Help Us Fix Security Bugs? *ArXiv abs/2112.02125* (2021).
- [Petroni et al.(2019)] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. 2019. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066* (2019).
- [Pialorsi and Russo(2007)] Paolo Pialorsi and Marco Russo. 2007. *Introducing microsoft® linq*. Microsoft Press.
- [Prenner et al.(2022)] Julian Aron Prenner, Hlib Babii, and Romain Robbes. 2022. Can OpenAI’s Codex Fix Bugs?: An evaluation on QuixBugs. In *2022 IEEE/ACM International Workshop on Automated Program Repair (APR)*. 69–75. <https://doi.org/10.1145/3524459.3527351>
- [Raffel et al.(2020)] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* 21, 140 (2020), 1–67.
- [Ren et al.(2020)] Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, M. Zhou, Ambrosio Blanco, and Shuai Ma. 2020. CodeBLEU: a Method for Automatic Evaluation of Code Synthesis. *ArXiv abs/2009.10297* (2020). <https://doi.org/10.48550/arXiv.2009.10297>
- [Song and Lu(2014)] Linhai Song and Shan Lu. 2014. Statistical debugging for real-world performance problems. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications*. 561–578. <https://doi.org/10.1145/2660193.2660234>
- [Tufano et al.(2020)] Michele Tufano, Dawn Drain, Alexey Svyatkovskiy, Shao Kun Deng, and Neel Sundaresan. 2020. Unit Test Case Generation with Transformers. *ArXiv abs/2009.05617* (2020). <https://doi.org/10.48550/arXiv.2009.05617>
- [Tufano et al.(2018)] Michele Tufano, Cody Watson, Gabriele Bavota, Massimiliano di Penta, Martin White, and Denys Poshyvanyk. 2018. An Empirical Investigation into Learning Bug-Fixing Patches in the Wild via Neural Machine Translation. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 832–837. <https://doi.org/10.1145/3238147.3240732>
- [Wei et al.(2022)] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of Thought Prompting Elicits Reasoning in Large Language Models. *ArXiv abs/2201.11903* (2022).
- [Weimer et al.(2009)] Westley Weimer, ThanhVu Nguyen, Claire Le Goues, and Stephanie Forrest. 2009. Automatically finding patches using genetic programming. In *2009 IEEE 31st International Conference on Software Engineering*. 364–374. <https://doi.org/10.1109/ICSE.2009.5070536>
- [Wen et al.(2018)] 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)*. 1–11. <https://doi.org/10.1145/3180155.3180233>
- [Wong et al.(2021)] Chu-Pan Wong, Priscila Santiesteban, Christian Kästner, and Claire Le Goues. 2021. VarFix: Balancing Edit Expressiveness and Search Effectiveness in Automated Program Repair. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 354–366. <https://doi.org/10.1145/3468264.3468600>
- [Xia and Zhang(2022)] Chunqiu Steven Xia and Lingming Zhang. 2022. Less Training, More Repairing Please: Revisiting Automated Program Repair via Zero-Shot Learning. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Singapore, Singapore) (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 959–971. <https://doi.org/10.1145/3540250.3549101>
- [Ye et al.(2021)] He Ye, Matias Martinez, Thomas Durieux, and Martin Monperrus. 2021. A comprehensive study of automatic program repair on the QuixBugs benchmark. *Journal of Systems and Software* 171 (2021), 110825.