



 Latest updates: <https://dl.acm.org/doi/10.1145/3696630.3728701>

RESEARCH-ARTICLE

## LLPut: Investigating Large Language Models for Bug Report-Based Input Generation

**ALIF AL HASAN**, Jahangirnagar University, Dhaka, Bangladesh

**SUBARNA SAHA**, Jahangirnagar University, Dhaka, Bangladesh

**MIA MOHAMMAD IMRAN**, Missouri University of Science and Technology, Rolla, MO, United States

**TARANNUM SHAILA ZAMAN**, University of Maryland, Baltimore County (UMBC), Baltimore, MD, United States

**Open Access Support** provided by:

**University of Maryland, Baltimore County (UMBC)**

**Missouri University of Science and Technology**

**Jahangirnagar University**



PDF Download  
3696630.3728701.pdf  
05 April 2026  
Total Citations: 0  
Total Downloads: 753

**Published:** 28 July 2025

**Citation in BibTeX format**

FSE Companion '25: 33rd ACM  
International Conference on the  
Foundations of Software Engineering  
June 23 - 28, 2025  
Trondheim, Norway

**Conference Sponsors:**  
SIGSOFT

# LLPut: Investigating Large Language Models for Bug Report-Based Input Generation

Alif Al Hasan  
Jahangirnagar University  
Dhaka, Bangladesh  
alif.stu2017@juniv.edu

Subarna Saha  
Jahangirnagar University  
Dhaka, Bangladesh  
subarna.stu2019@juniv.edu

Mia Mohammad Imran  
Missouri University of Science and Technology  
Rolla, Missouri, USA  
imranm@mst.edu

Tarannum Shaila Zaman  
University of Maryland Baltimore County  
Baltimore, Maryland, USA  
zamant@umbc.edu

## Abstract

Failure-inducing inputs play a crucial role in diagnosing and analyzing software bugs. Bug reports typically contain these inputs, which developers extract to facilitate debugging. Since bug reports are written in natural language, prior research has leveraged various Natural Language Processing (NLP) techniques for automated input extraction. With the advent of Large Language Models (LLMs), an important research question arises: how effectively can generative LLMs extract failure-inducing inputs from bug reports? In this paper, we propose LLPut, a technique to empirically evaluate the performance of three open-source generative LLMs—LLaMA, Qwen, and Qwen-Coder—in extracting relevant inputs from bug reports. We conduct an experimental evaluation on a dataset of 206 bug reports to assess the accuracy and effectiveness of these models. Our findings provide insights into the capabilities and limitations of generative LLMs in automated bug diagnosis.

## CCS Concepts

- **Software and its engineering** → **Software defect analysis.**

## Keywords

Bug Report, Empirical Analysis

### ACM Reference Format:

Alif Al Hasan, Subarna Saha, Mia Mohammad Imran, and Tarannum Shaila Zaman. 2025. LLPut: Investigating Large Language Models for Bug Report-Based Input Generation. In *33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)*, June 23–28, 2025, Trondheim, Norway. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3696630.3728701>

## 1 Introduction

Bug reports serve as the primary medium through which users and testers communicate software issues to developers [18].

These reports typically include a detailed natural language description of the problem, aiding developers in reproducing and diagnosing the bug. As a result, bug reports are a crucial component of software maintenance. Their content is not only essential for developers in maintaining software systems but also forms the foundation for automated tools that assist in the complex tasks of bug detection and resolution [18]. Whenever developers receive a bug report, the first step is to reproduce the bug [62, 64]. It is necessary to reproduce the bug to be confirmed and to observe the behavior of the bug. Bug reproduction also enables further analysis of the issue [10].

To reproduce a bug, developers need the failure-inducing input [60, 61]. These inputs are test cases or commands that cause a software failure during execution [29]. Developers also rely on inputs to diagnose and localize bugs. Ideally, bug reports clearly specify failure-inducing input commands. However, unstructured and ambiguous information in bug reports often makes them difficult to use [68]. Additionally, many bug reports describe inputs in natural language rather than technical terms. For example, instead of writing the command *chmod filename* to change a file's mode, a reporter might say, *change the file's permission mode*. This use of natural language makes it harder for new developers to extract the necessary inputs from bug reports.

Since bug reports are written in natural language, researchers have conducted numerous studies [38] using NLP [21] techniques to analyze bug reports and extract relevant information. This is the first study to systematically evaluate the effectiveness of LLMs in extracting input commands from bug reports. Recently, LLMs have assisted NLP researchers in various fields by analyzing human-written documents and extracting information from them [21].

The task of extracting *input commands* from bug report can be categorized as information extraction task. In recent years, the use of generative LLMs for information extraction has been popular [6, 19, 51, 58, 66], including software engineering tasks [37, 42]. They have been used in a zero-shot, one-shot, and few-shot manner to extract specific tasks. In software engineering, while LLMs have been used for tasks such as code summarization, bug localization, code generation, and vulnerability detection [3–5, 27, 31, 37, 39, 42, 59], extracting



This work is licensed under a Creative Commons Attribution 4.0 International License.

*FSE Companion '25, Trondheim, Norway*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1276-0/2025/06

<https://doi.org/10.1145/3696630.3728701>

input commands in bug reports has been not previously performed. This gap inspired us to investigate the effectiveness of the LLMs in extracting manual commands from bug reports. Therefore, we propose a technique LLPut to analyze LLM's performance in extracting inputs from bug report.




Figure 1: Overview of LLPut Methodology

Our proposed methodology, LLPut, consists of two main components: (i) dataset preparation and (ii) application of various LLM models for input extraction. Figure 1 provides an overview of LLPut's methodology. We begin by collecting and annotating bug reports to construct a structured dataset. As an initial baseline, we apply a BERT-based NLP technique to evaluate its effectiveness in extracting failure-inducing inputs. Subsequently, we employ three generative LLM models—LLaMA [56], Qwen [7], and Qwen-Coder [32]—to evaluate the performance of LLMs in extracting inputs from bug reports. Our dataset, codes, annotation instructions and prompts are available in this repo [1].

## 2 Dataset

To analyze the performance of the models in generating inputs from bug reports, we first collect and annotate a dataset of bug reports.

### 2.1 Dataset Collection

In this work, we focus on the Linux coreutils project, which includes 68 distinct utility applications [16] and offers a diverse range of functionalities. We collect bug reports from the online open source bug repository, *Red Hat Bugzilla - Bug List* [50]. We filter bug reports by selecting *Fedora* and *Red Hat Linux* as the product categories and *coreutils* as the target component. This filtering yields a total of 779 bug reports. We then retrieve these reports using the **REST API** and parse the data into two separate CSV files: one containing bug descriptions with their corresponding IDs and another containing IDs along with all other relevant attributes. In the future, we plan to expand our dataset by collecting data from various web-based applications and from other different open-source bug repositories.

To refine the dataset, we manually review the bug reports and exclude those that merely announce new *coreutils* releases. After this filtering process, we obtain a final dataset of 753 bug reports. Table 1 presents a sample from our dataset.

### 2.2 Dataset Annotation

Two authors manually annotate a randomly selected 250 reports from the refined dataset of 753 reports and discuss until they reach to consensus. They further discard reports with no

or vague descriptions (i.e., bug reports had nothing beyond an imprecise title, or any details that could help a user understand the bug). Afterward, we end up with a subset of 206 reports for further analysis. We give annotators the following instructions:

*Read the entire bug report thoroughly before beginning annotation. Pay attention to both explicit statements and implicit cues that might indicate the exact conditions which triggered the anomaly. For each bug report, determine the specific commands/instructions and in which order triggered that state.*

The detailed instruction is available in the project repository. The annotators carefully examine each report to identify the exact sequence of commands needed to reproduce the reported bugs.

Annotators note that many bug reports lack proper structure or omit essential details for understanding the issue. In these cases, annotators thoroughly read each report and extract the necessary information, including the expected outcome, observed outcome, and steps to reproduce the bug. Some reports only describe the encountered issue and how the user reached that state. The annotators analyze these cases, interpret the context, and reconstruct the steps to ensure accurate replication.

After annotation, We find that in 149 bug reports, input commands or test cases exist. 54 cases did not contain input commands or test cases. We mark them as 'None'. And in 3 cases annotators were uncertain as the actual commands were not explicit. We also mark them as 'None'.

## 3 Experimental Evaluation

In this research work, we evaluate LLPut, by considering two research questions:

**RQ 1:** How effective is BERT in accurately extracting user input commands from unstructured bug reports?

**RQ 2:** How effective is generative LLMs in extracting inputs from bug report?

### 3.1 Evaluation Metric

As *input commands* are needed to exactly match in order to run correctly, we compare the NLP (BERT model) extracted input commands and Generative LLM extracted input commands against human-annotated ones using BLEU score (Bilingual Evaluation Understudy) [45]. BLEU score is a standard metric in machine translation to evaluate text sequence similarity. In the software engineering domain, BLEU score is often used to evaluate the quality of generated bug reports, code comments, and code summarization.

**3.1.1 BLEU-score.** : The BLEU score is a standard metric used to assess the quality of machine-generated text by comparing it to human-written reference text [45]. It measures the similarity between the generated and reference texts by evaluating the overlap of n-grams. A higher BLEU score indicates a closer resemblance to the reference, making it particularly useful for

Table 1: Sample Data

id	title	Description	Input Command(s)
2296026	sort -V order differs from sort when there are no numerics	When sorting two lines with no numeric parts, sort -V differs from sort. [...] <b>How reproducible:</b> Always Note: The results are the same (as expected) if lines “a” and “abc” are sorted (i.e., without the “-xyz”): printf "a\nabc\n" sort -V   a abc [...]	1. printf "a-xyz\nabc-xyz\n" sort   2. printf "a-xyz\nabc-xyz\n" sort -V
2248503	uname -i doesn't work	coreutils-9.1-12.fc38.x86_64 [root@buildserver:~]\$ uname -i unknown This is breaking scripts expecting to get x86_64.	1. uname -i
749704	date cannot easily produce lower case formatted output	Suppose one wants lower case abbreviated day-of-week names. By default, %a produces capitalized names, e.g., Fri. Using %^a correctly produces upper case. Using %#a unfortunately produces the same as %^A, i.e., FRI. Quite clearly, FRI is not the opposite case of Fri. Instead of an 'opposite case' modifier [...]	1. date +%a 2. date +%^a 3. date +%#a 4. date +%%/a

tasks like machine translation and text generation. The BLEU score is computed using the following formula:

$$BLEU = BP \cdot \exp \left( \sum_{n=1}^N w_n \cdot \log(p_n) \right)$$

where,  $N$  represents the maximum n-gram order considered in the calculation;  $p_n$  is the *precision score* for n-grams of size  $n$ , reflecting the proportion of matching n-grams between the generated and reference texts;  $w_n$  is the weight assigned to each n-gram level, typically set to  $\frac{1}{N}$  to ensure equal importance across all n-gram orders; and  $BP$  denotes the *brevity penalty*, applied to discourage excessively short outputs. It equals 1 if the generated text is longer than the reference and less than 1 otherwise.

**3.1.2 BLEU Score Interpretation.** : BLEU score above 0.30 generally indicates that the generated text is understandable, while score exceeding 0.50 suggests high fluency and quality [20]. This scale has been widely applied in various fields [15, 36, 53], including software engineering [26, 33].

The BLEU score is sensitive to the chosen n-gram order. While a 4-gram model is commonly used in standard evaluations [20, 45], it may not always be suitable for all contexts. Since the extracted inputs tend to be brief (median is 4), BLEU-2 yields a more accurate reflection of text quality, consistent with approaches used in similar research [33]. Therefore, our analysis primarily relies on BLEU-2 for a more precise evaluation of the generated outputs.

### 3.2 RQ1: Effectiveness of NLP in Extracting Inputs

To answer this RQ, we adopt *token classification* approach [23]. In recent years, BERT based transformer models have been

widely used for token classification tasks [13, 55]. We investigate the effectiveness of BERT in extracting user input commands from unstructured bug reports.

**3.2.1 Dataset.** As our dataset has empty input commands, i.e., bug reports that did not contain any user input commands, we exclude them as they do not contribute to the command extraction task and may introduce noise into the training process. We end up with 149 data instances. The dataset was then split into training (80%) and testing (20%) sets.

**3.2.2 Model Configuration.** We use *bert-base-uncased* version. We fine-tune the model for a token classification task by utilizing the Huggingface transformers library [22]. The primary objective was to perform token classification specifically for extracting user input commands from unstructured bug reports. To achieve this, the model was trained to classify tokens into two categories: '1' (insider command) and '0' (outside a command). We use AdamW optimizer with a learning rate of  $2e-5$ . The model then trained over ten epochs with a batch size of 16.

**3.2.3 Results.** We observe that the BERT model performed very poorly. The BLEU-2 score in the test set greater than or equal to 0.5 was only once (3.33%) and in the 0.3-0.5 range only once again (3.33%), with no exact matches between predictions and ground truth. Manual inspection reveal that 40% of time the model often failed to capture a single token in the sequences, and rest 60% of the time either misclassified command tokens or partially extracted commands.

We hypothesize that the primary issue is the small size of our dataset, as tasks like command extraction typically require large-scale data [40, 54]. Additionally, the model struggled with complex command structures, such as multi-line commands or those embedded within descriptive text, leading to frequent errors.

These results highlight BERT’s limitations in handling unstructured bug reports with limited dataset and its inability to generalize to unseen command formats. To address these challenges, we explore *generative language models* [9], which offer greater flexibility and can operate in zero-shot and one-shot settings, enabling better generalization without extensive fine-tuning.

### 3.3 RQ2: Effectiveness of Generative LLMs in Extracting inputs

To answer this research question, We first design a prompt to extract the input commands from the bug reports. Afterwards, we use the same prompt for all models to extract input commands from bug reports.

**3.3.1 Experiment Setup.** We experiment with two distinct prompting strategies — zero-shot and one-shot — to evaluate the adaptability of LLMs to minimal-context tasks [17, 63, 69]. In the zero-shot setting, models are asked to complete tasks without any examples, relying solely on their pre-trained knowledge. In contrast, the one-shot approach provides a single illustrative example within the prompt to guide the model’s response. After preliminary testing and qualitative assessment, we determined that the one-shot prompt led to more structured and contextually accurate outputs. The finalized prompt design used in the experiments is detailed in section 3.3.4.

We then provide all models with the exact same prompt for each report and collect their responses alongside our manually gathered command sequences.

**3.3.2 Models Configuration.** We experiment the performance of three distinct open-source LLMs generative. We particularly focus on open-source models as they are more transparent, customizable, and accessible for research purposes. Below we briefly explain the models:

**LLaMA:** We utilized the *llama-3.3-70B* model, Meta AI’s latest iteration in the Llama series at the time of this experiment. With its 70 billion parameters, this model excels in both zero-shot and few-shot scenarios, demonstrating advanced reasoning skills and proficiency in handling knowledge-intensive tasks [56]. Its balanced performance across diverse domains makes it an ideal benchmark for general-purpose tasks.

**Qwen:** We employed the *qwen2.5:32b-instruct* version, a 32-billion parameter model developed by Alibaba’s AI research team. Optimized specifically for instruction-following tasks, Qwen demonstrates a strong ability to comprehend task-oriented prompts and generate structured outputs that align closely with user expectations [7]. Its design emphasizes prompt adherence and logical task decomposition, making it particularly suitable for procedural tasks [7].

**Qwen-coder:** We incorporated the *qwen2.5-coder:32b* variant, a specialized adaptation of the Qwen model fine-tuned for programming and code-related tasks [32]. This variant enhances the model’s capability in generating syntactically accurate and semantically coherent code snippets, making it

an ideal candidate for tasks involving software development, code analysis, and algorithmic reasoning [32].

**3.3.3 Parameters.** As generative LLM’s output may vary, we set the temperature 0, so that the output varies least while experimenting. We utilized Ollama platform for accessing these models [43].

**3.3.4 Prompt Design.** The structure of our input command extraction prompt is intended to mimic a real-world scenario where a user is going through bug reports and finding the exact command to replicate the bug report from the command line. The prompt is shown as below:

**Input Command Extractor Prompt**

I will provide you with a Linux coreutils bug description that was reported by a developer in Bugzilla. Your task is to determine the exact command(s) or test case(s) required to reproduce the bug. If a reproducible command or test case exists, write only the command or test case. If no reproducible command or test case is available, write None.

Example:  
**Input (Bug Description):** A user reports that the ls command incorrectly sorts filenames when using the -v option. The issue occurs when filenames contain both numbers and letters, leading to unexpected sorting behavior.  
touch file1 file2 file10 file20 ls -v  
**Output (Expected Response):** touch file1 file2 file10 file20 ls -v

Here is the reported Bug Description: < *insert Bug Description* >  
**Task:** Given the above bug description, identify the commands or test cases required to reproduce the bug.

**3.3.5 Results.** The range of BLEU scores and their respective count for the three different models using unigram, bigram, trigram, and four-gram are shown in Table 2. Note here that in some cases both the extracted command and human annotated command were empty, e.g., there were no commands in the bug report. The BLEU score is typically used for non-empty text comparisons. As, theoretically, two empty strings are treated as a perfect match, we consider the BLEU score in such cases 1.

From Table 2, we can observe that Qwen performed slightly better than LLaMA and Qwen-Coder. Considering BLEU-2, in 62.62% cases Qwen has a score  $\geq 0.5$ , while LLaMA and Qwen-Coder achieve 55.34% and 56.8%. In the case of score  $< 0.3$ , Qwen-Coder performs worst, with 38.35% scoring less than 0.3.

We observe that LLaMA has 48 cases, Qwen has 81 cases and Qwen-Coder has 67 cases where BLEU-2 score is exactly 1, i.e., the human annotation exactly matches with the model

**Table 2: BLEU Score Distribution Across Different Models.**

Model	Range	BLEU-1 Count (%)	BLEU-2 Count (%)	BLEU-3 Count (%)	BLEU-4 Count (%)
LLaMA	< 0.3	74 (35.92%)	76 (36.89%)	77 (37.38%)	80 (38.83%)
	0.3-<0.5	14 (6.8%)	16 (7.77%)	19 (9.22%)	20 (9.71%)
	≥ 0.5	118 (57.28%)	114 (55.34%)	110 (53.4%)	106 (51.46%)
Qwen	< 0.3	65 (31.55%)	66 (32.04%)	66 (32.04%)	69 (33.5%)
	0.3-<0.5	11 (5.34%)	11 (5.34%)	15 (7.28%)	13 (6.31%)
	≥ 0.5	<b>130</b> <b>(63.11%)</b>	<b>129</b> <b>(62.62%)</b>	<b>125</b> <b>(60.68%)</b>	<b>124</b> <b>(60.19%)</b>
Qwen-Coder	< 0.3	78 (37.86%)	79 (38.35%)	80 (38.83%)	82 (39.81%)
	0.3-<0.5	9 (4.37%)	10 (4.85%)	11 (5.34%)	11 (5.34%)
	≥ 0.5	119 (57.77%)	117 (56.8%)	115 (55.83%)	113 (54.85%)

extracted outputs. Out of these, in 38 cases all the models have exact matches.

In 81 cases all 3 models have BLEU-score greater than or equal to 0.5. In those 81 cases, the median of the input command to human annotation is 4, and the median of the input command generated by LLaMA, Qwen, and Qwen-Coder are respectively 5, 4, and 4. This is likely because when the models were able to extract correct commands, they were similar to human annotations. On the other hand, in 38 cases all 3 models have BLEU-score less than 0.3. In those 38 cases, the median to human annotation is 4, and the median to LLaMA, Qwen, and Qwen-Coder are respectively 18, 0, and 0. This indicates that when the models were incorrect, a) LLaMA likely extracted a command that were dissimilar to the human annotation, and b) Qwen and Qwen-coder likely predicted 'None' in majority of these cases.

**3.3.6 Error Analysis.** To understand the models' errors, we analyzed 38 instances where all three models had a BLEU score below 0.30. We focus particularly on the LLaMA model, as it produced fewer 'None' predictions, resulting in more samples with interpretable text.

Two authors carefully examined the human-annotated input commands and the LLM-generated commands. They employed thematic analysis to categorize the errors [52]. Specifically, One author first examined the commands and identified initial themes, which were then evaluated and refined by another author. They collaboratively discussed any inconsistencies and reached a consensus to finalize the categorization. Our analysis identifies three main error categories:

**Correct Extraction but Different Wordings:** There were 25 cases where the models generated slightly different variations of the input command, rather than the exact wordings. In those cases, while the models are technically correct, their generated command wordings were different.

**Failed Extraction:** There are 10 cases where the models failed to extract the input commands or test cases from the bug report where the annotators perceived otherwise.

**Subjectivity in Annotation:** There were 3 cases where it was not clear if the input commands or test cases were actually in

the bug report. Those cases where the model suggested that there were likely input commands, however, annotators were uncertain. Those are the cases, which further needed to be verified by the domain experts.

Table 3 shows the examples of these errors.

## 4 Discussion and Future Scope

Our error analysis reveals several avenues for enhancing the extraction of input commands from bug reports. Although our study relies primarily on BLEU scores, future work should integrate additional metrics—such as semantic similarity and context-aware measures—to capture subtle phrasing variations and provide a more comprehensive evaluation. Refining prompt design with dynamic, multi-stage approaches that incorporate contextual feedback may further align model outputs with human annotations.

Another promising direction is the integration of domain-specific knowledge. Our experiments indicate that models sometimes struggle with implicit or context-dependent commands; thus, fine-tuning on specialized datasets from diverse software domains could enhance their understanding of command patterns and improve extraction accuracy.

Finally, addressing the challenges of annotation subjectivity, where the presence of commands is ambiguous, may benefit from continuous expert feedback during training and evaluation. Broadening the evaluation framework to encompass bug reports from different domains and languages—and exploring hybrid approaches that blend rule-based heuristics with LLM outputs—will likely lead to more robust, context-aware command extraction systems and more effective debugging in real-world software development environments.

## 5 Related Work

Our work builds upon and extends previous research in two main areas: extracting structured information and how Prompt-based techniques for information extraction from bug reports. Below we discuss them.

**Extracting Structured Information from Bug Reports:** Extracting information from bug report has been a matured research area [8, 65]. Prior work has leveraged neural language techniques for bug report analysis, developing automated tools and frameworks that enhance bug identification and reduce manual effort significantly. Bettenburg identified which elements in bug reports are most valuable to developers [8]. Key advances include Rastkar et al.'s and Mani et al.'s work on automatic summarization of bug reports [41, 49]. Chaparro et al. developed methods for detecting missing information in bug descriptions [11, 12], while Imran et al. introduced techniques for finding missing information from bug reports [34]. Feng et al. analyzed secure device vulnerabilities through NLP-based techniques [25]. Zhu et al. proposed a graph-based neural model for automated bug localization with bug report decomposition and code hierarchical network [70]. Zhao et al. developed a tool that extract S2R from bug reports [67]. Ci-borowska et al. proposed a BERT-based technique for bug localization [14]. Ouedraogo et al. proposed an approach for

Table 3: Example of Errors

Error Type	Description	Input Command	LLaMA Output
Correct Extraction but Different Wordings	coreutils-9.3 brought changes to the behavior of the -v option which broke some of my automation scripts [...] Instead they only output this information with -debug [...]	1. cp -v -u <source> <destination> 2. mv -v -u <source> <destination>	1. cp -u -v file1 file2 2. mv -i -v file3 file4 3. cp -u -v -debug file1 file2 4. mv -i -v -debug file3 file4
Failed Extraction	(1) You mounted two partitions (2) You unmount the /mnt (formatted with ext3) (3) mount mpath2 (you unmounted in step 2) on /tmp/test/ (4) Now , you do df and expect this - [...]	1. mount /dev/mpath1 /tmp/test/umount /mnt 2. mount /dev/mpath2 /tmp/test/df	(None)
Subjectivity in Annotation	command install sets wrong permissions [...] Actual results: -rwxr-xr-x 1 root [...] Expected results: -rw-r--r-- 1 root [...]	(None)	1. bash install /home/vikram/Desktop/RPM-GPG-KEY-AL-CAPCOM /etc/pki/rpm-gpg/[...]

automatic test case generation after extracting relevant information from bug reports [44].

**Information Extraction from Bug Reports Through Prompting:** Recent advances in generative LLMs have introduced prompt-based techniques for bug report analysis. Feng et al. proposed a technique to automatically reproduce the bugs from bug reports through prompt engineering [24]. Plein did a feasibility study with LLMs on test case generation from bug reports [47]. Wen et al. analyzed static bug report warnings [57]. Ahmed et al. experimented on how capable GPT-4 and GPT-3.5 are to fix hardware based security bugs [2]. Kang et al. proposed a prompt based framework for bug reproduction [35]. Guan et al. introduced a knowledge-aware prompting framework that utilizes LLMs to generate test cases for model optimization bug [28]. Paul et al. applied zero-shot prompting technique using GPT-3.5-turbo for automatic program repair and found that they don't perform well [46].

Unlike those works, our empirical study specifically focus on how good *Open-source LLMs* are on extracting *input commands* from Linux Coreutils bug reports.

## 6 Threats to Validity

In this section, we discuss four types of threats, similar to prior research related to LLM[30, 48].

**Threats to Internal Validity:** Data contamination significantly threatens the internal validity of our study, as some of the bug reports we utilized may already be present in the training data of the evaluated LLMs.

**Threats to External Validity:** The study focuses on bug reports from Linux coreutils, specifically those reported on Redhat Bugzilla, which may limit the generalizability of our findings. While the Linux coreutils bug reports are diverse and complex, they represent only a subset of web application bug reports. Additionally, different websites may have distinct structures for bug reporting. In this study, we analyze bug reports from just one platform (Redhat), which could affect the performance of LLMs and limit the applicability of our conclusions to other domains of bug reports. To address this

limitation, we include a diverse set of 250 bug reports, spanning various categories and difficulty levels, to strengthen the robustness of our evaluation.

**Threats to Construct Validity:** We test the models in a one-shot setting, without additional fine-tuning or iterative interactions. This approach may not fully utilize the models' capabilities, as fine-tuning or interactive prompting could enhance performance on specific tasks. To address this limitation, we document all experimental procedures and make the scripts used for data preprocessing and solution evaluation publicly available, ensuring reproducibility and transparency.

**Threats to Conclusion Validity:** Our evaluation relies on the data annotation phase and the prompt used to assess the performance of the LLM. To minimize bias, we ensure that all evaluations are conducted using the same prompt. Additionally, we verify the manual data annotation by having two different authors independently review the annotations.

## 7 Conclusion

Our investigation of LLMs demonstrates that generative LLMs can effectively extract failure-inducing inputs from bug reports without extensive training data compared to BERT-based model, which performed poorly. The generative models showed promising results, with Qwen achieving the highest accuracy (62.62% outputs with BLEU-2  $\geq$  0.5). Key challenges identified include command wording variations, extraction failures, and annotation ambiguity. These findings suggest significant potential for integrating LLMs into bug reproduction workflows, reducing manual effort for developers. Future work should focus on domain-specific knowledge integration, refined prompting strategies, and hybrid approaches to further improve extraction accuracy and reliability across diverse software ecosystems.

## 8 Acknowledgments

This work was supported in part by NSF grants CCF-2348277 and CCF-2518445.

## References

- [1] 2025. *LLPut: Investigating Large Language Models for Bug Report-Based Input Generation*. Zenodo. doi:10.5281/zenodo.15092886
- [2] Baleegh Ahmad, Shailja Thakur, Benjamin Tan, Ramesh Karri, and Hammond Pearce. 2024. On hardware security bug code fixes by prompting large language models. *IEEE Transactions on Information Forensics and Security* (2024).
- [3] Toufique Ahmed and Premkumar Devanbu. 2022. Few-shot training llms for project-specific code-summarization. In *Proceedings of the 37th IEEE/ACM international conference on automated software engineering*. 1–5.
- [4] Toufique Ahmed, Kunal Suresh Pai, Premkumar Devanbu, and Earl Barr. 2024. Automatic semantic augmentation of language model prompts (for code summarization). In *Proceedings of the IEEE/ACM 46th international conference on software engineering*. 1–13.
- [5] Vishwanath Akuthota, Raghunandan Kasula, Sabiha Tasnim Sumona, Masud Mohiuddin, Md Tanzim Reza, and Md Mizanur Rahman. 2023. Vulnerability detection and monitoring using llm. In *2023 IEEE 9th International Women in Engineering (WIE) Conference on Electrical and Computer Engineering (WIECON-ECE)*. IEEE, 309–314.
- [6] Marwah Alaofi, Negar Arabzadeh, Charles LA Clarke, and Mark Sanderson. 2024. Generative information retrieval evaluation. In *Information Access in the Era of Generative AI*. Springer, 135–159.
- [7] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609* (2023).
- [8] Nicolas Bettenburg, Rahul Premraj, Thomas Zimmermann, and Sunghun Kim. 2008. Extracting structural information from bug reports. In *Proceedings of the 2008 international working conference on Mining software repositories*. 27–30.
- [9] Charlotte Blease, John Torous, Brian McMillan, Maria Hägglund, and Kenneth D Mandl. 2024. Generative language models and open notes: exploring the promise and limitations. *JMIR medical education* 10 (2024), e51183.
- [10] Patrick Joseph Casey. 1992. *Computer Program Debugging: An Engaging Problem-solving Environment*. Ph. D. Dissertation. New York, NY, USA. UMI Order No. GAX92-10525.
- [11] Oscar Chaparro, Carlos Bernal-Cárdenas, Jing Lu, Kevin Moran, Andrian Marcus, Massimiliano Di Penta, Denys Poshyvanyk, and Vincent Ng. 2019. Assessing the quality of the steps to reproduce in bug reports. In *Proceedings of the 2019 27th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*. 86–96.
- [12] Oscar Chaparro, Jing Lu, Fiorella Zampetti, Laura Moreno, Massimiliano Di Penta, Andrian Marcus, Gabriele Bavota, and Vincent Ng. 2017. Detecting missing information in bug descriptions. In *Proceedings of the 2017 11th joint meeting on foundations of software engineering*. 396–407.
- [13] Gunjan Chhablani, Abheesht Sharma, Harshit Pandey, Yash Bhartiya, and Shan Suthaharan. 2021. NLRG at SemEval-2021 Task 5: Toxic Spans Detection Leveraging BERT-based Token Classification and Span Prediction Techniques. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*. 233–242.
- [14] Agnieszka Ciborowska and Kostadin Damevski. 2022. Fast changeset-based bug localization with bert. In *Proceedings of the 44th international conference on software engineering*. 946–957.
- [15] Google Cloud. 2025. AutoML Evaluation | Cloud Translation | Google Cloud. <https://cloud.google.com/translate/docs/advanced/automl-evaluate>
- [16] coreutils [n. d.]. *Decoded: GNU coreutils*. Retrieved Feb. 24, 2025 from <https://www.maizure.org/projects/decoded-gnu-coreutils/>
- [17] Hai Dang, Lukas Mecke, Florian Lehmann, Sven Goller, and Daniel Buschek. 2022. How to prompt? Opportunities and challenges of zero-and few-shot learning for human-AI interaction in creative applications of generative models. *arXiv preprint arXiv:2209.01390* (2022).
- [18] Steven Davies and Marc Roper. 2014. What’s in a bug report?. In *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (Torino, Italy) (ESEM ’14)*. Association for Computing Machinery, New York, NY, USA, Article 26, 10 pages. doi:10.1145/2652524.2652541
- [19] Shumin Deng, Yubo Ma, Ningyu Zhang, Yixin Cao, and Bryan Hooi. 2024. Information extraction in low-resource scenarios: Survey and perspective. In *2024 IEEE International Conference on Knowledge Graph (ICKG)*. IEEE, 33–49.
- [20] Michael Denkowski and Alon Lavie. 2010. Choosing the right evaluation for machine translation: an examination of annotator and automatic metric performance on human judgment tasks. In *Proceedings of the 9th Conference of the Association for Machine Translation in the Americas: Research Papers*.
- [21] Jianshu Du, Yibo Wang, Wenting Zhao, Zhongfen Deng, Shuaiqi Liu, Renze Lou, Henry Zou, Pranav Narayanan Venkit, Nan Zhang, Mukund Srinath, et al. 2024. LLMs Assist NLP Researchers: Critique Paper (Meta-) Reviewing. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 5081–5099.
- [22] Hugging Face. 2025. Hugging Face. <https://huggingface.co>
- [23] Hugging Face. 2025. Token Classification - Hugging Face. <https://huggingface.co/tasks/token-classification>
- [24] Sidong Feng and Chunyang Chen. 2024. Prompting is all you need: Automated android bug replay with large language models. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*. 1–13.
- [25] Xuan Feng, Xiaojing Liao, X Wang, Haining Wang, Qiang Li, Kai Yang, Hongsong Zhu, and Limin Sun. 2019. Understanding and securing device vulnerabilities through automated bug report analysis. In *SEC’19: Proceedings of the 28th USENIX Conference on Security Symposium*.
- [26] Zhipeng Gao, Xin Xia, John Grundy, David Lo, and Yuan-Fang Li. 2020. Generating question titles for stack overflow from mined code snippets. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 29, 4 (2020), 1–37.
- [27] Qiuhan Gu. 2023. Llm-based code generation method for golang compiler testing. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 2201–2203.
- [28] Hao Guan, Guangdong Bai, and Yepang Liu. 2024. Large language models can connect the dots: Exploring model optimization bugs with domain knowledge-aware prompts. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 1579–1591.
- [29] Ralf Hildebrandt and Andreas Zeller. 2000. Simplifying failure-inducing input. *SIGSOFT Softw. Eng. Notes* 25, 5 (Aug. 2000), 135–145. doi:10.1145/347636.348938
- [30] Md Sifat Hossain, Anika Tabassum, Md. Fahim Arefin, and Tarannum Shaila Zaman. 2025. LLM-ProS: Analyzing Large Language Models’ Performance in Competitive Problem Solving. *arXiv:2502.04355 [cs.CL]* <https://arxiv.org/abs/2502.04355>
- [31] Soneya Binta Hossain, Nan Jiang, Qiang Zhou, Xiaopeng Li, Wen-Hao Chiang, Yingjun Lyu, Hoan Nguyen, and Omer Tripp. 2024. A deep dive into large language models for automated bug localization and repair. *Proceedings of the ACM on Software Engineering* 1, FSE (2024), 1471–1493.
- [32] Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [33] Mia Mohammad Imran, Preetha Chatterjee, and Kostadin Damevski. 2024. Uncovering the causes of emotions in software developer communication using zero-shot llms. In *Proceedings of the IEEE/ACM 46th international conference on software engineering*. 1–13.
- [34] Mia Mohammad Imran, Agnieszka Ciborowska, and Kostadin Damevski. 2021. Automatically selecting follow-up questions for deficient bug reports. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 167–178.
- [35] Sungrim Kang, Juyeon Yoon, and Shin Yoo. 2023. Large language models are few-shot testers: Exploring llm-based general bug reproduction. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2312–2323.
- [36] Alon Lavie. 2011. Evaluating the Output of Machine Translation Systems. In *Proceedings of Machine Translation Summit XIII: Tutorial Abstracts*. Xiamen, China. <https://aclanthology.org/2011.mtsummit-tutorials.3/>
- [37] Peng Li, Tianxiang Sun, Qiong Tang, Hang Yan, Yuanbin Wu, Xuan-Jing Huang, and Xipeng Qiu. 2023. CodeIE: Large Code Generation Models are Better Few-Shot Information Extractors. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 15339–15353.
- [38] Lutfun Nahar Lota, Tarannum Shaila Zaman, Mirza Mohammad Azwad, Labiba Farah, Abrar Chowdhury, Zaarin Anjum, Chadni Islam, and Abu Raihan Mostofa Kamal. [n. d.]. Recent Trends and Challenges in Using Nlp Techniques in Software Debugging: A Systematic Literature Review. Available at SSRN 5060080 ([n. d.]).
- [39] Guilong Lu, Xiaolin Ju, Xiang Chen, Wenlong Pei, and Zhilong Cai. 2024. GRACE: Empowering LLM-based software vulnerability detection with graph structure and in-context learning. *Journal of Systems and Software* 212 (2024), 112031.
- [40] Kai Ma, Yongjian Tan, Miao Tian, Xuejing Xie, Qinjun Qiu, Sanfeng Li, and Xin Wang. 2022. Extraction of temporal information from social media messages using the BERT model. *Earth Science Informatics* 15, 1 (2022), 573–584.
- [41] Senthil Mani, Rose Catherine, Vibha Singhal Sinha, and Avinava Dubey. 2012. Ausum: approach for unsupervised bug report summarization. In *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*. 1–11.

- [42] Ying Mo, Jiahao Liu, Jian Yang, Qifan Wang, Shun Zhang, Jingang Wang, and Zhoujun Li. 2024. C-ICL: Contrastive In-context Learning for Information Extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. 10099–10114.
- [43] ollama.com. 2025. Ollama. <https://ollama.com/>
- [44] Wendkūni C Ouédraogo, Laura Plein, Kader Kaboré, Andrew Habib, Jacques Klein, David Lo, and Tegawendé F Bissyandé. 2024. Extracting Relevant Test Inputs from Bug Reports for Automatic Test Case Generation. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*. 406–407.
- [45] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [46] Rishov Paul, Md Mohib Hossain, Mohammed Latif Siddiq, Masum Hasan, Anindya Iqbal, and Joanna Santos. 2023. Enhancing automated program repair through fine-tuning and prompt engineering. *arXiv preprint arXiv:2304.07840* (2023).
- [47] Laura Plein, Wendkūni C Ouédraogo, Jacques Klein, and Tegawendé F Bissyandé. 2024. Automatic generation of test cases based on bug reports: a feasibility study with large language models. In *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*. 360–361.
- [48] Shireesh Reddy Pyreddy and Tarannum Shaila Zaman. 2025. EmoXpt: Analyzing Emotional Variances in Human Comments and LLM-Generated Responses. *arXiv:2501.06597* [cs.LG] <https://arxiv.org/abs/2501.06597>
- [49] Sarah Rastkar, Gail C Murphy, and Gabriel Murray. 2014. Automatic summarization of bug reports. *IEEE Transactions on Software Engineering* 40, 4 (2014), 366–380.
- [50] redhat\_bug [n. d.]. *Red Hat Bugzilla – Main Page*. Retrieved June 01, 2020 from <https://bugzilla.redhat.com/>
- [51] Oscar Sainz, Iker García-Ferrero, Rodrigo Agerri, Oier Lopez de Lacalle, German Rigau, and Eneko Agirre. 2024. GoLLIE: Annotation Guidelines improve Zero-Shot Information-Extraction. In *ICLR*.
- [52] B. Saunders, J. Sim, T. Kingstone, S. Baker, J. Waterfield, B. Bartlam, H. Burroughs, and C. Jinks. 2018. Saturation in qualitative research: exploring its conceptualization and operationalization. *Quality & quantity* 52 (2018).
- [53] Sanja Seljan, Marija Brkić, and Tomislav Vičić. 2012. BLEU Evaluation of Machine-Translated English-Croatian Legislation. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*. Istanbul, Turkey, 2143–2148.
- [54] Hyeon Jin Shin, Jeong Yeon Park, Dae Bum Yuk, and Jae Sung Lee. 2020. BERT-based spatial information extraction. In *Proceedings of the Third International Workshop on Spatial Language Understanding*. 10–17.
- [55] Wenjun Sun, Hanh Thi Hong Tran, Carlos-Emiliano González-Gallardo, Mickaël Coustaty, and Antoine Doucet. 2024. LIT: Label-Informed Transformers on Token-Based Classification. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 144–158.
- [56] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [57] Cheng Wen, Yuandao Cai, Bin Zhang, Jie Su, Zhiwu Xu, Dugang Liu, Shengchao Qin, Zhong Ming, and Tian Cong. 2024. Automatically inspecting thousands of static bug warnings with large language model: How far are we? *ACM Transactions on Knowledge Discovery from Data* 18, 7 (2024), 1–34.
- [58] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, Yang Wang, and Enhong Chen. 2024. Large language models for generative information extraction: A survey. *Frontiers of Computer Science* 18, 6 (2024), 186357.
- [59] Aidan ZH Yang, Claire Le Goues, Ruben Martins, and Vincent Hellendoorn. 2024. Large language models for test-free fault localization. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*. 1–12.
- [60] Tingting Yu, Tarannum S. Zaman, and Chao Wang. 2017. DESCRy: reproducing system-level concurrency failures. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (Paderborn, Germany) (ESEC/FSE 2017)*. Association for Computing Machinery, New York, NY, USA, 694–704. doi:10.1145/3106237.3106266
- [61] Tarannum Shaila Zaman and Tariqul Islam. 2022. ReDPro: an automated technique to detect and regenerate process-level concurrency failures. In *Proceedings of the 2022 ACM Southeast Conference (Virtual Event) (ACMSE '22)*. Association for Computing Machinery, New York, NY, USA, 106–112. doi:10.1145/3476883.3520207
- [62] Tarannum Shaila Zaman, Macharla Hemanth Kishan, and Lutfun Nahar Lota. 2025. Gender Dynamics in Software Engineering: Insights from Research on Concurrency Bug Reproduction. *arXiv:2502.20289* [cs.SE] <https://arxiv.org/abs/2502.20289>
- [63] J Diego Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–21.
- [64] Andreas Zeller. October 2005. *Why Programs Fail: A Guide to Systematic Debugging*.
- [65] Jie Zhang, Xiaoyin Wang, Dan Hao, Bing Xie, Lu Zhang, and Hong Mei. 2015. A survey on bug-report analysis. *Sci. China Inf. Sci.* 58, 2 (2015), 1–24.
- [66] Zikang Zhang, Wangjie You, Tianci Wu, Xinrui Wang, Juntao Li, and Min Zhang. 2025. A Survey of Generative Information Extraction. In *Proceedings of the 31st International Conference on Computational Linguistics*. 4840–4870.
- [67] Yu Zhao, Kye Miller, Tingting Yu, Wei Zheng, and Minchao Pu. 2019. Automatically extracting bug reproducing steps from android bug reports. In *Reuse in the Big Data Era: 18th International Conference on Software and Systems Reuse, ICSR 2019, Cincinnati, OH, USA, June 26–28, 2019, Proceedings 18*. Springer, 100–111.
- [68] Yu Zhao, Tingting Yu, Ting Su, Yang Liu, Wei Zheng, Jingzhi Zhang, and William G.J. Halfond. 2019. ReCDroid: Automatically Reproducing Android Application Crashes from Bug Reports. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 128–139. doi:10.1109/ICSE.2019.00030
- [69] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. 2024. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *International Journal of Machine Learning and Cybernetics* (2024), 1–65.
- [70] Ziyue Zhu, Hanghang Tong, Yu Wang, and Yun Li. 2022. Enhancing bug localization with bug report decomposition and code hierarchical network. *Knowledge-Based Systems* 248 (2022), 108741.