



 Latest updates: <https://dl.acm.org/doi/10.1145/3788866>

RESEARCH-ARTICLE

Exploring Large Language Models in Resolving Environment-Related Crash Bugs: Localizing and Repairing

XUEYING DU, Fudan University, Shanghai, China

MINGWEI LIU, Sun Yat-Sen University, Guangzhou, Guangdong, China

HANLIN WANG, Fudan University, Shanghai, China

JUNTAO LI, Fudan University, Shanghai, China

XIN PENG, Fudan University, Shanghai, China

YILING LOU, Fudan University, Shanghai, China

Open Access Support provided by:

Fudan University

Sun Yat-Sen University



PDF Download
3788866.pdf
12 February 2026
Total Citations: 0
Total Downloads: 47

Accepted: 16 December 2025
Revised: 10 November 2025
Received: 14 April 2025

[Citation in BibTeX format](#)

Exploring Large Language Models in Resolving Environment-Related Crash Bugs: Localizing and Repairing

XUEYING DU, Department of Computer Science, Fudan University, China

MINGWEI LIU*, School of Software Engineering, Sun Yat-sen University, China

HANLIN WANG, Department of Computer Science, Fudan University, China

JUNTAO LI, Department of Computer Science, Fudan University, China

XIN PENG, Department of Computer Science, Fudan University, China

YILING LOU, Department of Computer Science, Fudan University, China

Software crash bugs cause unexpected program behaviors or even abrupt termination, thus demanding immediate resolution. However, resolving crash bugs can be challenging due to their complex root causes, which can originate from issues in the source code or external factors like third-party library dependencies. Large language models (LLMs) have shown promise in software engineering tasks, leveraging their extensive pre-training on text and code corpora. However, existing research predominantly focuses on the capability of LLMs to localize and repair code-related crash bugs, leaving their effectiveness in resolving environment-related crash bugs in real-world software unexplored.

To fill this gap, we conducted the first comprehensive study to assess the capability of LLMs in resolving real-world environment-related crash bugs, using a newly constructed dataset of 100 representative crash bugs. We first systematically compare LLMs' performance in resolving code-related and environment-related crash bugs with varying levels of crash contextual information. Our findings reveal that the LLM performs better in resolving code-related crash bugs compared to environment-related ones. Specifically, localization is the primary challenge for resolving code-related crashes, while repair poses a greater challenge for environment-related crashes. Furthermore, we investigate the impact of different prompt strategies on improving the resolution of environment-related crash bugs, incorporating different prompt templates and multi-round interactions. Building on this, we further explore an advanced active inquiry prompting strategy, which leverages the self-planning capabilities of LLMs to conduct systematic and continuous questioning aimed at identifying potential environmental factors that contribute to crashes. Based on these explorations, we propose IntDiagSolver, an interactive methodology designed to enable precise crash bug resolution through ongoing engagement with LLMs. Extensive evaluations of IntDiagSolver on a dataset of 41 crash bugs across multiple LLMs (including GPT-3.5, GPT-4, Claude, CodeLlama, DeepSeek-R1, and Qwen-3-Coder) demonstrate consistent improvements in resolution accuracy, with substantial enhancements ranging from 9.1% to 43.3% in localization and 9.1% to 53.3% in repair. Furthermore, the strong performance of IntDiagSolver on the latest expanded multilingual dataset of 42 crash bugs highlights its strong generalizability and effectiveness on previously unseen data.

*Corresponding authors (liumw26@mail.sysu.edu.cn)

Authors' Contact Information: Xueying Du, xueyingdu21@m.fudan.edu.cn, Department of Computer Science, Fudan University, Shanghai, China; Mingwei Liu, liumw26@mail.sysu.edu.cn, School of Software Engineering, Sun Yat-sen University, Zhuhai, China; Hanlin Wang, wanghanlin23@m.fudan.edu.cn, Department of Computer Science, Fudan University, Shanghai, China; Juntao Li, 22210240197@m.fudan.edu.cn, Department of Computer Science, Fudan University, Shanghai, China; Xin Peng, pengxin@fudan.edu.cn, Department of Computer Science, Fudan University, Shanghai, China; Yiling Lou, yilinglou@fudan.edu.cn, Department of Computer Science, Fudan University, Shanghai, China.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 1557-7392/2026/1-ART

<https://doi.org/10.1145/3788866>

CCS Concepts: • **Software and its engineering** → *Maintaining software; Application specific development environments*; • **Computing methodologies** → *Artificial intelligence*.

Additional Key Words and Phrases: Large Language Model, Crash Bug Resolution, Bug Localization

1 Introduction

Crash bugs pose significant challenges, potentially leading to program termination and disrupting user operations [28, 35]. These failures typically originate from two distinct categories: code-related defects (*i.e.*, bugs in the source code, such as null pointer dereferences) and external environmental factors (*e.g.*, improper OS versions, or missing third-party dependencies) [71].

Although extensive efforts have been made in automatically localizing and repairing code-related crash bugs[22–24, 56, 77, 78], addressing environment-related crash bugs remains challenging. This complexity arises from the multifaceted nature of environmental dependencies, including system configuration discrepancies, version incompatibilities, permission constraints, and hardware-specific requirements. Moreover, the numerous environmental factors involved within a software system make achieving an end-to-end solution more challenging. Existing techniques often rely on adopting similar solutions from online forums like Stack Overflow (SO) [52, 54, 71], which lack generalizability, particularly for previously undocumented issues. Our preliminary analysis of 100 crash-related Stack Overflow threads reveals that 68% of root causes stem from environmental factors, underscoring the prevalence and practical significance of this issue.

Recent advances in large language models (LLMs) demonstrate remarkable capabilities in code comprehension [17, 72] and generation tasks [30, 60, 86]. While some existing researches have explored LLMs' potential in bug fixing[65, 77], these efforts predominantly focus on code-related crash bugs. The effectiveness of LLMs in resolving environment-related crash bugs in real-world software remains unexplored.

To address these knowledge gaps, we conducted the first comprehensive study to assess the capability of LLMs in resolving real-world environment-related crash bugs. Our investigation focused on evaluating LLMs' effectiveness in both localizing and repairing environment-related crash bugs using various prompts and interaction methods. Specifically, we: (i) systematically compared the performance of LLMs in resolving both code-related and environment-related crash bugs under different levels of contextual information, (ii) investigated how various prompt strategies on improving the resolution of environment-related crash bugs, incorporating different prompt templates and multi-round interactions, and (iii) further explored the effectiveness of a novel active inquiry prompting strategy that leverages self-planning to enhance multi-round prompting. We aim to answer the following two research questions:

- **RQ1 (Comparative Effectiveness):** How do LLMs perform in localizing and repairing code-related versus environment-related crash bugs with varying levels of contextual information?
- **RQ2 (Enhancing Resolution with Prompt Engineering):** How do existing prompt strategies, as well as our proposed advanced prompt strategy, improve the effectiveness of LLMs in resolving environment-related crash bugs?
 - **RQ2.a (Influence of Prompt Templates):** How do different prompt templates affect LLMs' ability to resolve crash bugs?
 - **RQ2.b (Impact of Multi-Round Prompts):** What is the impact of employing multi-round prompts on the effectiveness of LLMs in resolving crash bugs?
 - **RQ2.c (Impact of Active Inquiry Prompt with Self-Planning):** What is the impact of employing the advanced active inquiry prompt on the effectiveness of LLMs in resolving crash bugs?

Our study of 100 real-world crash bugs reveals key findings. Firstly, the LLM demonstrates more excellent proficiency in resolving *code-related crash bugs* than *environment-related crash bugs*. Secondly, we identify fault localization as the primary challenge in LLMs' resolution of code-related crash bugs, while repair poses a greater

challenge for environment-related crash bugs. Thirdly, continuous interaction with LLM proves to be effective in acquiring relevant knowledge and enhancing the effectiveness of environment-related crash bug resolution. Moreover, by stimulating LLMs' proactive questioning ability for self-guided planning, we observe a significant improvement in resolution as the LLM navigates the resolution process through step-by-step localization.

Based on our findings, we introduce a methodology IntDiagSolver to enhance the interaction with LLMs for more accurate resolution of crash bugs. While our primary focus is on environment-related crash bugs, we ensure the applicability of our approach to code-related crash bugs by proposing distinct prompt templates highlighting different aspects of crash contexts for both types of crash bugs respectively, coupled with diverse multi-round interaction strategies. Additionally, we devise a strategy enabling LLMs to guide the repair process, which is particularly beneficial for users with limited crash bug knowledge. We then evaluate our proposed methodology addressing the following research questions:

- **RQ3 (Effectiveness of IntDiagSolver):** How does IntDiagSolver enhance the ability of LLMs to effectively resolve crash bugs?
- **RQ4 (Generalizability of IntDiagSolver):** How well does IntDiagSolver generalize to different scenarios?
 - **RQ4.a (Generalizability across LLMs):** To what extent can the effectiveness of IntDiagSolver in resolving crash bugs be generalized across different LLMs?
 - **RQ4.b (Generalizability across Programming Languages):** To what extent can the effectiveness of IntDiagSolver in resolving crash bugs be generalized across different programming languages?

The results demonstrate that IntDiagSolver consistently improves the accuracy of resolving both code-related and environment-related crash bugs across six different LLMs, with substantial enhancements ranging from 9.1% to 43.3% in localization and 9.1% to 53.3% in repair. Notably, there is a remarkable improvement from 0/30 to 16/30 in the number of successfully repaired environment-related crash bugs on GPT-3.5. Furthermore, results on the latest expanded multilingual environment-related crash bug dataset further highlights the strong generalizability and effectiveness of IntDiagSolver on previously unseen data, with localization improvements of 12.0% to 16.6% and repair improvements of 19.0% to 21.4% across two advanced LLMs.

The contributions of this work are summarized as follows:

- *First Comprehensive Study:* Represents the first comprehensive investigation into LLMs' ability to localize and repair environment-related crash bugs. Our benchmark and results are available at online replication package [16].
- *Insightful Findings:* Delves into the strengths and limitations of crash bug resolution using LLMs, providing a comparative analysis of LLM performance in code-related and environment-related crash bugs localization and repair. Furthermore, we explore the impact of different prompt enhancement techniques and interaction methods on resolution outcomes.
- *Innovative Methodology:* Introduces the novel methodology IntDiagSolver that enhances the interaction with LLMs for more accurate crash bug resolution.

2 RQ1: Comparative Effectiveness

In this RQ, we perform a comparative analysis of the LLM's ability to localize and repair code-related and environment-related crash bugs, considering varying levels of crash contextual information.

Code-related crash bugs refer to those caused by defects within the software's source code. Repairing such bugs typically involves modifying the source code itself. In contrast, **environment-related crash bugs** arise from issues or incompatibilities in external factors required for the software's execution. These may include the operating system version or configuration, dependency library versions, hardware resources, network status, or the availability and configuration of third-party services. Typical resolutions for environment-related bugs

involve adjusting the relevant environmental configurations, such as modifying configuration files or updating dependencies, without the need to change the source code.

2.1 Benchmark

Previous researches [61, 66, 77] exploring the potential of LLMs in bug resolution typically utilize benchmarks such as QuixBugs [47], Defects4j [43], ManySSuBs4J [44], and UnifiedBugDataset [33]. However, these benchmarks primarily focus on bugs originating from faulty source code, neglecting those arising from environmental issues, which are prevalent and sometimes more challenging for developers [69, 89]. Furthermore, the widely used QuixBugs benchmark, which includes only 40 bugs in relatively simple scenarios, such as basic algorithms, fails to capture the complexity of real-world crash bugs.

To comprehensively compare the effectiveness of LLMs in addressing code-related and environment-related crash bugs in real-world applications, we create a new benchmark derived from SO threads, incorporating both code-related and environment-related issues. SO provides rich contextual information on buggy code's purposes and dependencies.

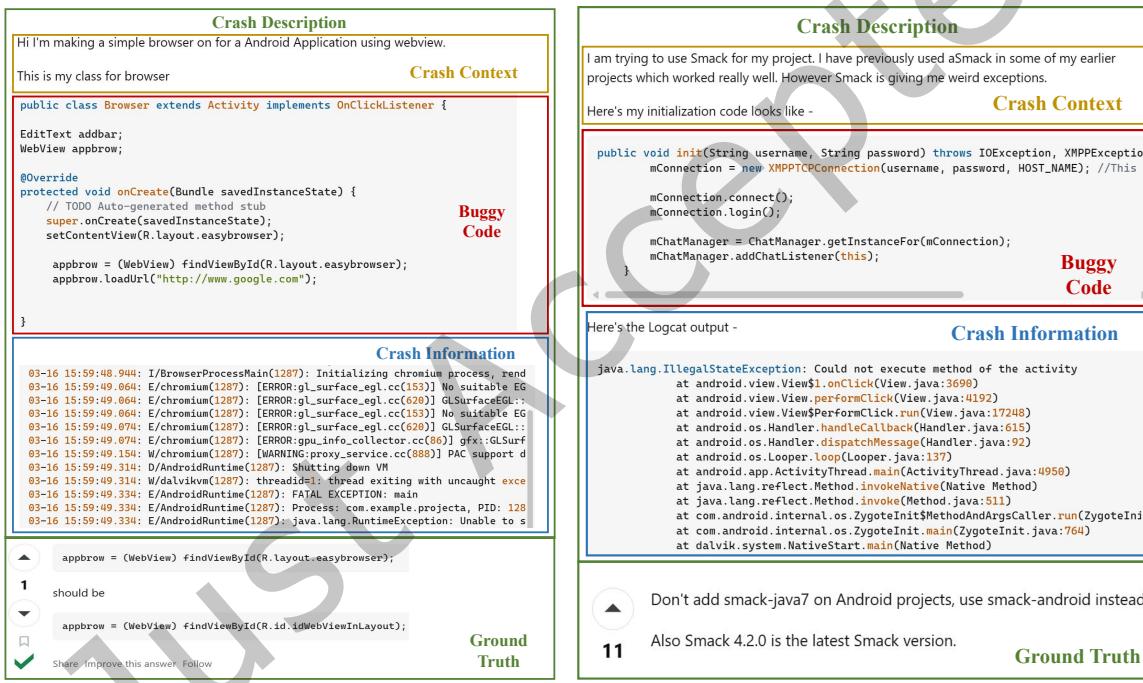


Fig. 1. Examples of Crash Bugs

2.1.1 Benchmark Format. We extracted crash descriptions (Cra-Des) for each crash bug from the question sections of SO threads. These descriptions facilitate a comprehensive analysis of LLMs' effectiveness in localizing and repairing crashes within various information contexts. Specifically, crash descriptions consist three primary parts of information:

- **Buggy Code (B-Code)** : This snippet of code triggers the crash and often contains syntax or semantic errors that lead to unexpected program behavior.

- **Crash Information (Cra-Info)** : This segment includes the exception type, error message, and crash trace. Exception type identifies the error, *e.g.*, divide-by-zero. The error message provides a human-readable description of the error. The crash trace offers a stack trace of the program execution at the time of the crash, assisting in locating the error.
- **Crash Context (Cra-Cont)** : This section provides a natural language description of the broader crash context, including symptoms, potential causes, relevant inputs, the buggy code’s purpose, configurations, or dependencies possibly contributing to the error. Such details are crucial for replicating the crash and grasping its underlying cause.

We obtained the ground truth for localizing and repairing each crash bug from the accepted answers in the corresponding SO thread. Figure 1(a) illustrates examples of both code-related and environment-related crash bugs from our benchmark, specifically extracted from SO threads #22442280 and #30322026.

2.1.2 Benchmark Construction Procedure. We automated the construction of a pool of SO threads enriched with detailed crash descriptions, from which we selected 100 samples possessing high-quality ground truths to serve as our dataset. In alignment with previous research [35, 71], we focus on crash bugs in the Java program due to their prevalence [52].

Establishing Data Pool. To establish a benchmark encompassing sufficient contextual information for both code-related and environment-related crash localization and repair, we selected high-quality threads from the SO data dumps [1] based on the following inclusion criteria: (1) Thread titles or tags containing the keyword “Java”, (2) Thread titles or tags containing the keywords “exception” or “error”, (3) Threads with an accepted answer, (4) Threads with at least one positive vote for their question, (5) Threads with concrete error symptoms (*i.e.*, code or stack trace) to ensure sufficient context for understanding the crash bug, and (6) Threads containing at least one common Java exception type (*e.g.*, `NullPointerException`). Furthermore, to collect a comprehensive list of common Java exception types for Criterion-6, we systematically parsed 35,773 Java libraries from Maven Central [3] based on the Libraries.io dataset [2] and JDK 1.8 [6]. In this way, we collected 67,248 SO threads related to crash bugs as the pool.

High-Quality Sample Selection. We initially randomly sampled 100 threads from the curated pool, and then two authors independently assessed the accepted answers in each thread, which are the answers verified for correctness by the original question askers. Following prior work[31, 53, 55, 70], we treated these accepted answers as the ground truth solutions for the corresponding crash bugs. Only threads with ground truth solutions consistently rated as high quality by both assessors were included in the final benchmark. High-quality solutions were defined as those providing detailed code modifications or problem-solving steps, in contrast to merely offering vague troubleshooting tips or general fixing directions without specific guidance. Additionally, we filtered out cases that could not be unambiguously categorized as either code-related or environment-related. We continued sampling and assessment until we had 50 qualified SO threads for both code-related and environment-related types, *i.e.*, 100 crash bugs in total. Since the cases were randomly selected, the difficulty levels of the code-related and environment-related cases are expected to follow a natural probability distribution. Table 1 summarizes the statistics of our dataset.

Table 1. Dataset Statistics. The “Buggy code lines” column shows the average number of buggy code lines per instance, while the “Crash trace length” and “Crash description length” indicate the average token counts of the crash trace and the natural language description in the corresponding thread, respectively.

Type	Buggy Code Lines	Crash Trace Length	Crash Description Length
Code-related	47.24	783	418
Environment-related	34.57	2773	406

2.2 Setup and Metrics

We manually reviewed LLM’s responses and conducted separate qualitative analyses to assess its localization and repair capabilities. We opted for the gpt-3.5-turbo-1106 model for this RQ. This choice was driven by its more rapid response and cost-effectiveness [67].

Correctness of Localization and Repair. *Code-related Crash Bugs.* In the localization evaluation, a solution was deemed correct if the line of code it identified for modification matched the faulty code line identified in the benchmark. In the repair evaluation, the LLM was required to provide an accurate code patch that matched the one in the benchmark to be considered correct. *Environment-related Crash Bugs.* In the localization evaluation, the LLM’s localization was considered correct if it identified the precise cause of the exception (e.g., a library version problem or a missing dependency problem). In the repair evaluation, a specific repair plan was required (e.g., upgrading or downgrading to a specific library version).

The evaluation was conducted independently by two authors. In cases of disagreement, a third evaluation was carried out using the same prompt to determine the correctness of the response, following a majority-win strategy. Consequently, the Kappa coefficient [57] for code-related crash bugs localization and repair was 0.89 and 0.92 respectively; for both localization and repair in environment-related crash bugs was 0.97, indicating almost perfect agreement. It’s worth noting that only a small percentage of queries (5%, 5 out of 100) required this third evaluation to reach a consensus.

Solution Number. Typically, the LLM produces a single accurate solution with high confidence in its diagnosis; With reduced confidence, it may propose multiple potential solutions. Consequently, we recorded the number of solutions the LLM provided for each crash bug and manually verified the correctness of localization and repair for each solution.

Accuracy & Success. We calculated the accuracy rates for localization and repair, reflecting the proportion of correct solutions for localization/repair to the total solutions offered. The formula for calculating accuracy is as follows:

$$\text{Accuracy} = \frac{1}{N} \left(\sum_{i=1}^N \frac{\text{CorrectSol.Number}_i}{\text{TotalSol.Number}_i} \right) \times 100\% \quad (1)$$

A crash bug is deemed successfully localized/repairs if any one of the solutions within a response is identified as correct.

2.3 Resolving Code-related Crash Bugs

2.3.1 Design. We first assess the LLM’s ability to localize and repair 50 code-related crash bugs in our benchmark, utilizing the aforementioned three types of crash description information (i.e., Buggy Code, Crash Information, and Crash Context).

Prompt Design. A series of prompt templates (Basic-Prompt-1 to Basic-Prompt-4) were designed for interacting with the LLM, covering different levels of granularity of information.

Basic-Prompt-1: This is my code: [B-CODE] I’m getting an exception, how do I fix it?

Basic-Prompt-2: This is my code: [B-CODE] I’m getting [CRA-INFO], how do I fix it?

Basic-Prompt-3: [CRA-CONT] This is my code: [B-CODE] I’m getting [CRA-INFO], how do I fix it?

Basic-Prompt-4: [CRA-DES] This is my error lines: [LOC] how to fix it?

Interaction Procedure. For each code-related crash bug in our benchmark, we conduct experiments with Basic-Prompt-1 to 3 to evaluate the LLM’s effectiveness with different levels of provided information. Basic-Prompt-1 includes only the buggy code, Basic-Prompt-2 adds the buggy code along with additional crash information, and Basic-Prompt-3 further incorporates crash context. For crash bugs that could not be correctly resolved with any of the aforementioned prompts, Basic-Prompt-4 was employed. This prompt specifically pinpointed the *localization result in the buggy code lines (LOC)*, identifying the exact lines requiring modification. This simplified

the task for the LLM, allowing it to focus solely on repairing the crash bug without localization, thereby enabling an examination of the model’s repair capabilities in isolation.

Table 2. Experimental Results on Crash Bug Resolution. The “Information” column contains input data, while the “Localization” and “Repair” columns show the number of successfully localized and repaired crash bugs. “Loc. with Multi.” indicates cases where the LLM localized the bug but offered multiple solutions.

Type	Prompt	Information	Localization	Repair	Loc. with Multi.	Localization Accuracy	Repair Accuracy
Code related	Basic-Prompt-1	B-Code	20/50	20/50	2	0.380	0.380
	Basic-Prompt-2	B-Code + Cra-Info	30/50 (+10)	28/50 (+8)	2	0.580 (+52.6%)	0.540 (+42.1%)
	Basic-Prompt-3	B-Code + Cra-Info + Cra-Cont	42/50 (+22)	39/50 (+19)	3	0.807 (+112.3%)	0.747 (+96.6%)
	Basic-Prompt-4	B-Code + Cra-Info + Cra-Cont + LOC	46/50 (+26)	43/50 (+23)	3	0.887 (+133.3%)	0.827 (+117.5%)
Env. related	Basic-Prompt-2	B-Code + Cra-Info	7/50	5/50	5	0.063	0.047
	Basic Prompt 3	B-Code + Cra-Info + Cra-Cont	32/50 (+25)	20/50 (+15)	21	0.341 (+441.3%)	0.193 (+310.7%)

2.3.2 *Results.* Table 2 illustrates the experimental results of the LLM’s ability to resolve code-related crash bugs. The multifaceted analysis yields the following main conclusions (Finding 1-4).

Response Analysis. We collected the responses of the LLM for different prompts and manually analyzed the contents of the responses. For code-related crash bugs, the LLM’s responses can encompass several aspects, including confirming code functionality, explaining error messages, identifying potential causes, offering textual fixes, presenting repaired code, discussing alternatives, and providing warnings.

Overall Resolution Capability. The experimental results presented in Table 2 indicate that the LLM has significant potential in fixing code-related crash bugs. Specifically, when provided with full crash description information, the LLM was able to correctly localize 42 out of 50 crash bugs (84.0%) and repaired 39 of them (78.0%). Furthermore, the LLM provided a unique and clear answer for the majority of repaired crash bugs, with only 7.9% (3 out of 38) repaired crash bugs providing multiple possible solutions.

Finding 1: The LLM excels in resolving code-related crash bugs, displaying proficiency in bug localization and repair.

Contextual Information Impact. The LLM demonstrates the ability to resolve code-related bugs using only the information available within the buggy code. As shown in Table 2, the LLM successfully localized and repaired 40% (20 out of 50) of crash bugs when provided with only the buggy code using Basic-Prompt-1. This indicates that the LLM can effectively understand and resolve code-related bugs to a certain extent, relying solely on the information within the buggy code.

Moreover, Table 2 illustrates that providing more detailed crash description information significantly enhanced the LLM’s performance in localizing and repairing crash bugs. **Notably, both the number of localized and repaired bugs increased, indicating that a more effective crash description improved the LLM’s bug localization accuracy.**

The significance of exception type and error message in crash localization is highlighted by the responses of the LLM and the analysis of two specific cases: crash bugs 68199510 [7] and 65084069 [8] (the SO thread IDs of the crash bugs). In both cases, the LLM provided environment-specific solutions, leading to misdiagnosis. This is because the exception types *NoClassDefFoundError* and *NoSuchMethodError* are commonly associated with environmental issues. Our investigation further reveals that complete stack traces in code-related crash bugs play a relatively minor role compared to the exception type information. This finding was validated through additional experiments on specific crash cases. In crash bug 22928450 [9], the LLM successfully generated the

correct repair solution when provided with only the buggy code, crash context, and exception type. However, when the exception type information was replaced with a stack trace that lacked explicit exception type details, the LLM failed to accurately localize the crash bug. Therefore, **to improve repair effectiveness, it is beneficial to explicitly emphasize exception type information in natural language within the crash description.**

Comprehensive crash context information also enhances the LLM’s ability to localize issues, encompassing various aspects such as the intended functionality of the program, detailed symptoms of the crash (e.g., specific conditions under which the crash occurs versus when the program operates normally, user-generated test cases, and previous unsuccessful repair attempts). Furthermore, it is worth noting that we observed no cases where a crash bug successfully resolved with limited information became unresolved when richer context was introduced. This indicates that providing more detailed context consistently improve bug resolution effectiveness.

Finding 2: Enhanced crash description information significantly improves the LLM’s ability to localize code-related bugs. Exception type and error message are crucial for precise localization, and additional context in natural language descriptions further enhances its performance.

Localization vs. Repair. Among the queries complete crash descriptions, only 7.1% (3 out of 42) of correctly localized bugs remained unresolved. Moreover, in cases where the LLM initially failed to locate the crash bug, providing the actual buggy code line enabled successful repair in 4 out of 7 instances. This indicate that the LLM’s resolving bottleneck mainly lies in localization. Once the LLM correctly localize the crash bugs, it demonstrates a high success rate in successfully repair. From the results shown in Table 2, only 8.0% (4 of 50) crash bugs could not be repaired despite being able to localize the fault line.

Finding 3: The primary bottleneck in resolving code-related crash bugs lies in localization rather than repair.

Code Length Impact. To investigate the effect of code length on crash bug resolution, we compared cases with fewer than 50 lines of code to those with more than 50 lines, using the complete crash context (basic-prompt-3). For crash bugs involving complex buggy code (exceeding 50 lines), the localization and repair accuracies are 11/14 (78.6%) and 10/14 (71.4%), respectively. For crash bugs with simpler buggy code (fewer than 50 lines), the localization and repair accuracies are 31/36 (86.1%) and 29/36 (80.6%). Furthermore, only 2 out of 14 (14.3%) cases with more than 50 lines of buggy code could be correctly localized when only the buggy code was provided. These results indicate that longer code significantly impairs both localization and repair performance.

Finding 4: Longer code leads to a poorer localization and repair ability, necessitating additional crash description information for effective resolution.

2.4 Resolving Environment-related Crash Bugs

2.4.1 Design. Mirroring the design for code-related crash bugs (refer to Section 2.3.1), we further explore the LLM’s capability in addressing environment-related crash bugs with varying crash description information. Given that not all environmental crash descriptions include code snippets and that the root cause of environment-related crash bugs may not be related to code implementation, we classify environmental crashes based on natural language descriptions and non-natural language information(*i.e.*, code snippets and crash information). For each case, we first experimented with Basic-Prompt-2 and Basic-Prompt-3. If accurate repair could not be achieved with these prompts, we then employed Basic-Prompt-4 for further investigation.

2.4.2 Results. The LLM exhibits diminished proficiency in environment related crash bugs resolution than code-related crash bugs, as evidenced by the lower localization and repair accuracy and more candidate solutions provided, indicating an inherent uncertainty in the LLM’s bug localization. Based on the results presented in Table 2 and our detailed analysis of specific examples, we draw several conclusions (Finding 5-9) from different aspects.

Overall Resolution Capability. Table 2 shows that the LLM successfully identify 42 code-related crash bugs and 32 environment-related crash bugs with full crash description. Moreover, among the accurately localized responses, 21 out of 32 (65.6%) cannot provide a single clear solution. In some cases (e.g. crash bug 28097042 [10] and 24305296 [11]), the LLM gives five or more possible solutions for the same crash bug, resulting in a low accuracy. Different from code-related crash bugs, full stack trace is more crucial for environment-related crash bugs diagnosis, as which contains a great deal of environment-related information.

Finding 5: The LLM shows lower proficiency in resolving environment-related crash bugs compared to code-related ones due to its inability to pinpoint the root cause, resulting in multiple imprecise candidate solutions.

Localization vs. Repair. Differing from the observations made while resolving code-related crash bugs, 12 instances with full crash descriptions were successfully localized but remained unresolved (only 3 instances in resolving code-related crash bugs). This is due to the LLM can more easily leverage error messages and crash traces to accurately localize the type of environment-related crash bugs (such as permission, version, or IP address issues), but it often struggles to provide specific repair steps to resolve these issues. Therefore, guiding the LLM to generate explicit and actionable repair strategies remains a key challenge for future research.

Finding 6: The LLM lacks specificity in providing solutions for environment-related crash bugs.

Contextual Information Impact. Table 2 shows that providing buggy code and crash information successfully identify 60% (30 of 50) of code-related crash bugs. Adding crash context only results in the successful location of 12 more bugs. Addressing environment-related crash bugs through buggy code and crash information identifies only 14% (7 of 50) of such bugs, with only 10% (5 of 50) successfully repaired. However, the inclusion of crash context significantly increases success rates to 64% (32 of 50) for localization and 40% (20 of 50) for fixing environment-related crash bugs. This highlights the importance of crash context for environment-related crash bugs repair, compared to code-related crash bugs. However, leveraging the LLM to resolve environment-related crash bugs is challenging due to the subjective nature of crash context and its dependence on natural language, requiring more expertise from the user.

Finding 7: In contrast to code-related crash bugs, the repair for environment-related crash bugs relies heavily on crash context.

We also observe that the LLM is unable to reason and localize the causes of a bug when the relevant environmental information is not included in the crash description. For instance, in the case of crash bug 49871007 [12], the exception description and stack trace only mention “Resteasy 3.1.4.Final”. Yet, the actual library version that needed modification was “javax.ws.rs-api”, while “Resteasy 3.1.4.Final” is an implementation of the JAX-RS-API 2.1 specification. In such cases, the LLM fails to effectively reason and localize the cause.

Finding 8: The LLM is unable to reason and localize the causes of crash bugs when the relevant environmental information is not included in the crash description.

Considering the high frequency and impact of version problems in environment-related crash bugs, which represent a significant proportion (19 out of 50 in our benchmark), we conducted a specialized analysis of these issues. We observed that the LLM typically recommends upgrading to the latest version to fix library version-related crash bugs, e.g., bug 24305296 [11]. However, for cases where upgrading or downgrading to a specific version of the library was the solution (e.g., bugs 43320334 [13] and 37771758 [14]), the LLM couldn’t provide the right version number.

Finding 9: The LLM typically recommends upgrading to the latest version to solve library version-related issues.

3 RQ2: Enhancing Resolution with Advanced Prompts

In this section, we explore advanced prompts to enhance the LLM's ability to address environment-related crash bugs. To this end, we conduct experiments that investigate two key aspects of advanced prompts: prompt templates and multi-round prompts. Specifically, we first improve the design of prompt templates for single-round interactions (as detailed in Section 3.3), and then we enhance the resolution performance by facilitating continuous interactions in multi-round scenarios (as discussed in Section 3.4 and 3.5). This study includes a comprehensive analysis of a user study comprising 10 distinct cases. The specific details of these cases are provided in Section 3.1, while the evaluation metrics are explained in Section 3.2.

3.1 Study Cases

We meticulously selected 10 diverse cases that remained unresolved despite complete crash descriptions provided in RQ1. These cases span a range of issues including dependency version conflicts, configuration settings errors, network issues, and user authority issues. To ensure the representativeness of the selection, cases were sampled proportionally, ensuring at least one instance from each category of issues. Detailed information for each case is presented in Table 3. These 10 cases will be used for specific experiments and analysis concerning each sub-problem.

Table 3. Summary of Crash Bugs Used as Case Studies

Bug ID	Crash Bug Issue	Solution Summary
30322026	version issues	replace smack-java with smack-android
51370703	version issues	downgrade the version of SLF4J from 1.8.0 to 1.7.X
37771758	version issues	downgrade the version of Glassfish server from 4.1.1 to 4.1.0
60210757	version issues	upgrade the version of mongo-kafka-connect later than 1.0.0
49871007	version issues	upgrade the JAX-RS-API to 2.1
43320334	version issues	upgrade Chrome, Chrome driver and Selenium
48637658	network issues	replace the local IP address with the corresponding IP address for the Android emulator
67044715	authority issues	make sure the user has access to the table
42933291	configuration settings issues	install one of the normal arphic fonts or the texlive variant.
39858254	configuration settings issues	add the TrustStore configuration to server.xml

3.2 Setup and Metrics

We conducted the same manual examination to evaluate the correctness of the LLM's responses in terms of localization and repair as outlined in RQ1. To mitigate the influence of the LLM's inherent randomness, we **conducted three trials for each case**. Therefore, the accuracy metric was calculated based on the proportion of successful localizations and repairs across all 30 trials. We opted for the gpt-3.5-turbo-1106 model in alignment with RQ1. Additionally, we rated the answers for usefulness, conciseness, and interactivity using a 4-point Likert scale [46] (1-disagree; 2-somewhat disagree; 3-somewhat agree; 4-agree) based on predefined statements (see the following). For each query, all scoring results were jointly determined by two participants through discussions.

- **Usefulness:** The LLM can provide useful solutions for resolving crash bugs, including a detailed explanation of the root cause of the issue and a comprehensive solution.
- **Conciseness:** The LLM can provide clear and accurate solutions without redundant information.
- **Interactivity:** The conversation flows well across multiple rounds, without losing sight of previous context.

3.3 RQ2.a: Influence of Prompt Templates

Previous research has demonstrated that prompt templates have a significant impact on the LLM’s responses [75]. Therefore, we further explore how to improve prompt template design to enhance the LLM’s performance in repairing environment-related crash bugs. This section focuses on how to design high-quality prompt templates and introduces corresponding experimental designs and results.

3.3.1 Design. Building on our experimentation experience and drawing inspiration from [4], we craft three distinct prompt templates for environment-related crash bugs.

Multi-Solution-Prompt: Please show me all potential solutions.

Role-Play-Prompt: I want you to act as a fault localization and program repair expert. You will be able to provide detailed solutions to fix the given program crash.

Chain-of-Thought-Prompt: Please fix it step by step.

Firstly, we explore whether the LLM can produce the correct solution that was previously unattainable, by promoting the generation of all possible solutions through the utilization of **Multi-Solution-Prompt**.

Inspired by [29, 74, 88], the role-play prompt and chain-of-thought prompt may significantly enhance the LLM’s efficacy in localization and repair. Concatenating the above prompt with **Basic-Prompt-3**, we used these prompts to conduct comparative experiments on 10 studied cases. To mitigate any potential concerns with the randomness of the LLM, we ran three requests for each query and independently evaluated each response.

Table 4. Experimental Results of Exploring Advance Prompts

Prompt	Rounds	Localization	Repaired	Solution Num.	Usefulness	Conciseness	Interactivity
Basic	Single	14/30	3/30	5.8	1.6	1.5	/
Multi-Solution	Single	21/30 (+7)	9/30 (+6)	7.67	2.43 (+0.83)	0.56 (-0.94)	/
Role-Play	Single	23/30 (+9)	10/30 (+7)	5.23	2.83 (+1.23)	1.87 (+0.37)	/
Chain of Thoughts	Single	21/30 (+7)	6/30 (+3)	9.13	1.9 (+0.3)	0.2 (-1.3)	/
Basic + Multi-round	Multi.	25/30 (+11)	16/30 (+13)	4.77	2.93 (+1.33)	2.17 (+0.67)	3.37
Role-Play + Multi-round	Multi.	27/30 (+13)	18/30 (+15)	4.83	3.27 (+1.77)	2.2 (+0.7)	3.63

3.3.2 Results. The experimental results of various prompt templates are presented in Table 4. Compared to the basic prompt, the Multi-Solution-Prompt enables the LLM to generate a broader range of solutions, increasing the likelihood of finding a correct one, but resulting in reducing conciseness. Similarly, the Chain-of-Thought-Prompt introduces substantial redundancy, leading to a low conciseness score of 0.2, without significantly improving the LLM’s ability to derive a more accurate end-to-end solution. Compared to the Multi-Solution and Chain-of-Thought Prompts, which enhance localization and repair effectiveness but sacrifice conciseness, the Role-Play-Prompt excels in single-round interactions. It outperforms the basic prompt with 0.23 higher accuracy in repairs, a 1.23 higher usefulness score, and improved conciseness.

Furthermore, the Role-Play-Prompt improves the LLM’s ability to identify the root causes of crash bugs. For instance, while the first two prompts recognized issues related to the Smack version in crash bug 30322026, they failed to identify the mismatch between the Smack version and the Android project as the underlying cause, which the responses of Role-Play-Prompt successfully pinpointed. Thus, we conclude:

Finding 10: Role-play prompts significantly improve the LLM's ability to localize and repair crash bugs.

3.4 RQ2.b: Impact of Multi-round Prompts

We investigate the impact of multi-round interactions on improving the LLM's ability to resolve crash bugs, hypothesizing that iterative interactions can refine its responses and enhance overall performance.

3.4.1 Design. To evaluate this hypothesis, we conduct experiments assessing the effectiveness of providing feedback on the LLM's responses across two different scenarios.

- The initial answer did not accurately localize the root cause of the crash bugs (e.g., crash bug 37771758).
- The initial answer identified the root cause of the crash bugs, but the solution provided was too general and did not offer specific resolution steps (e.g., crash bugs 51370703 and 30322026).

For the first scenario, we devise the New-Solution-Prompt to request new solutions; for the second scenario, we use the Refinement-Prompt for the specific details of the solution. During our attempts, we found that Refinement-Prompt can obtain solution details for most issues. However, for issues related to library versions, Refinement-Prompt was insufficient in obtaining specific version numbers that needed to be upgraded or downgraded. Therefore, we designed specific prompts for such problems that require targeted questioning. For each crash bug, we employed both Basic-Prompt and Role-Play-Prompt in our experiments.

New-Solution-Prompt: I have tried the solution above, but the issue remains. Can you give me some new solutions?

Refinement-Prompt: Please provide more detailed information regarding the solutions mentioned above.

Version-Prompt: Which version of [Library 1] is compatible with my project/ [Library 2] version?

3.4.2 Results. Based on the specific analysis of the study cases and the results shown in Table 4, we can draw the following conclusions (Findings 11-13), each with a detailed analysis.

Table 4 illustrates that continuous interaction enhances the model's performance in both localization and repair accuracy, as well as in scoring for usefulness and conciseness, compared to single-round interactions. Notably, even with the basic prompt in continuous interaction, the results surpass those of the most effective single-round role-play prompt. This underscores the effectiveness of continuous interaction.

Finding 11: Continuously interacting with the LLM can help to further improve its ability to localize and repair crash bugs.

The experiments in Table 4 show that utilizing the Role-Play-Prompt for multi-turn interactions outperforms the basic prompt across all metrics. This suggests that the Role-Play-Prompt helps the LLM better understand the context and generate more accurate responses. Additionally, it significantly enhances the interaction score from 3.37 to 3.63. This improvement can be attributed to the Role-Play-Prompt aiding the LLM in maintaining contextual consistency and avoiding the repetition of previous solutions.

Finding 12: The use of Role-Play-Prompt significantly improves the LLM's ability to maintain contextual coherence and reduces instances of forgetting previous context, resulting in better resolving and interaction performance.

Based on our trials with crash bug 51370703, we provide the following conclusion as a supplement to finding 8.

Finding 13: the LLM appears to have knowledge about library versions that are discussed in official documentation. However, it requires targeted questioning to activate the corresponding knowledge.

For instance, during continuous interaction, a vague inquiry such as “*Which version of SLF4J should I use?*” lead the LLM to merely suggest upgrading to the latest available version or recommend version 1.8, as referenced in the bug report. In contrast, posing a more targeted question such as “*Which version of SLF4J is compatible*

with Logback 1.2.3 in my project?" can prompt the LLM to retrieve specific knowledge that *Logback version 1.2.3 explicitly requires the slf4j-api version 1.7.x*.

3.5 RQ2.c: Enhancing Multi-round Prompts with Self-planning

From the analysis above, we observe that even with advanced prompts such as the Role-Play Prompt and Chain-of-Thought prompt, the LLM struggles to accurately identify the root cause of a crash bug in a single attempt, limiting its ability to provide targeted solutions. As discussed in Section 3.4 and Finding 13, while the LLM possesses the necessary knowledge to address many crash bugs, effectively utilizing this knowledge requires precise strategies. Continuous interaction and targeted questioning rely on the questioner's ability to discern the true cause from multiple potential solutions, imposing high demands on the user. Furthermore, the lack of sufficient contextual information from the questioner constrains the LLM's ability to generate targeted solutions (Finding 8). Additionally, questioners may struggle to identify the relevant aspects of the broader environment-related context necessary for accurately diagnosing the crash bug.

To address these challenges, we empower the LLM to actively investigate potential crash-causing environmental factors through continuous proactive inquiry, guiding the resolution process step by step. By leveraging the LLM's self-planning capability [39], we propose an **active inquiry prompt** that guides the LLM in constructing an investigative plan. This plan examines various environmental factors that may contribute to crashes, prioritizes them based on diagnostic relevance, and acquires necessary information by posing targeted questions to the user. Although iterative feedback-based prompting has been explored in prior software engineering research [20, 41], the key innovation of the active inquiry prompt lies in its ability to stimulate the LLM to ask proactive and targeted questions. Through this iterative questioning process, the LLM progressively refines its understanding, gradually narrows down the root cause, and ultimately provides precise solutions. Moreover, this approach proves particularly beneficial for novice users who may lack specialized expertise, enabling them to more effectively resolve environment-related crash bugs.

3.5.1 Design and Results. We used the role-play prompt to enable the LLM for active inquiry. We refined the questioning method to avoid overly general queries and prevent overwhelming developers with multiple questions simultaneously. This included specifying certain question types and providing an example to guide the LLM's questioning toward more targeted and actionable suggestions. The resulting prompt stimulates the LLM's proactive questions, facilitating self-planning.

ActiveQ-Prompt: [Role-Play-Prompt] Additionally, please use the Socratic method of questioning to aid in accurate diagnosis. Note that the information in the question should be as specific as possible. For example, ...[AskOneQ-Prompt]

AskOneQ-Prompt: Note that only ask one question at a time, starting with the question you think is the most important.

By using this method, the LLM is able to ask one question at a time based on its priority, leading developers to targeted solutions. Utilizing the crash bug numbered 30322026 as an illustrative case, the LLM posed the following inquiry with the utilization of ActiveQ-Prompt and AskOneQ-Prompt:

My first question is: What is the version of the Smack library that you are using in this project?

After providing the corresponding information, partial the LLM's response would be as follows:

To fix this issue, you can try using a Smack library version that is specifically designed for Android, such as smack-android-extensions:4.4.4 or smack-android:4.4.4.

Based on the response, it is evident that the LLM accurately pinpointed the underlying cause of the bug, which pertains to the incompatibility of the employed smack-java7:4.1.0 library with the Android environment. Furthermore, the LLM suggested a correct library version as the detailed solution.

Finding 14: Employing active inquiry prompts to stimulate proactive questions for self-planning enhances the LLM's ability to proactively guide the repair process. This aids in identifying the root causes of crash bugs, proving particularly beneficial for novice users.

4 Methodology: IntDiagSolver

Due to the inherent non-determinism of environment-related crash bugs, their symptoms manifest through complex interactions between software components and runtime configurations, presenting significant challenges in developing end-to-end solutions. Based on our experimental findings, we propose IntDiagSolver, a novel approach for addressing environment-related crash bugs through continuous interaction with LLMs.

Figure 2 illustrates the overview of our approach. The proposed framework operates through a continuous interaction cycle between the user and the LLM, comprising three main phases: **Phase 1: Contextual Information Completion via LLM Self-Planning**. Given the bug report, this phase utilizes the LLM's self-planning capabilities to automatically generate and iteratively refine queries. These queries aim to gather essential contextual information that enriches the original bug report. **Phase 2: Solution Generation**. Utilizing the enriched context from the previous phase and the initial bug report, the LLM generates potential solutions. This generation process involves iterative refinement prompts designed to progressively yield a specific and actionable detailed solution. **Phase 3: Solution Validation**. In the final phase, we evaluate the correctness of the detailed solution through execution validation. If the execution succeeds, the verified correct solution is returned directly to the user. Conversely, if the execution fails, the LLM iteratively refines and generates a revised solution.

Figure 3 illustrates the detailed workflow with specific prompts utilized by IntDiagSolver. To enhance the generalizability of IntDiagSolver, we extend its application to effectively handle both environment-related and code-related crash bugs.

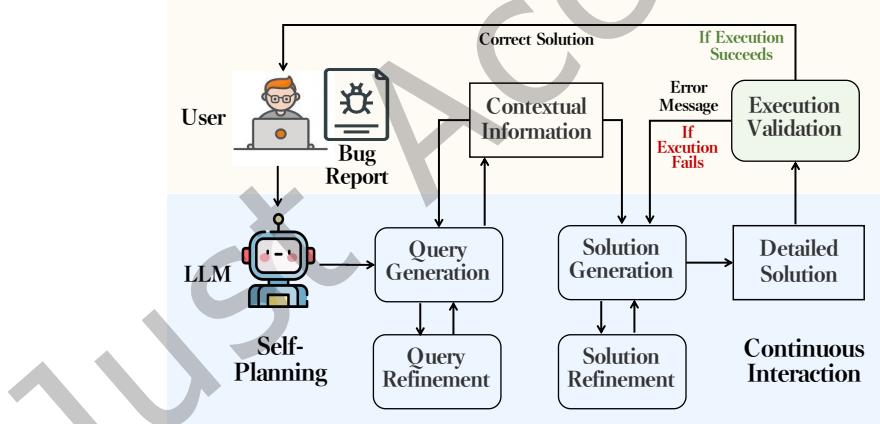


Fig. 2. Overview of IntDiagSolver

Concept Explanations. **Prompt provided by the users** is in a yellow box with crash details and instructions. The details of different prompts can be referred to the prompt templates listed in RQ1 and RQ2. **Response provided by LLMs** is in a blue box and can be categorized into questions and solutions. Questions may be multiple, but highlighting the question requirement will help the LLM generate a specific question (see section 3.5). Solutions can be rough or detailed (including optimized solution) with specific steps to fix the crash bugs. **Validation** is the repair attempt result in green box. **Decision** in diamond-shaped box indicating a selection, which based on manual judgment currently.

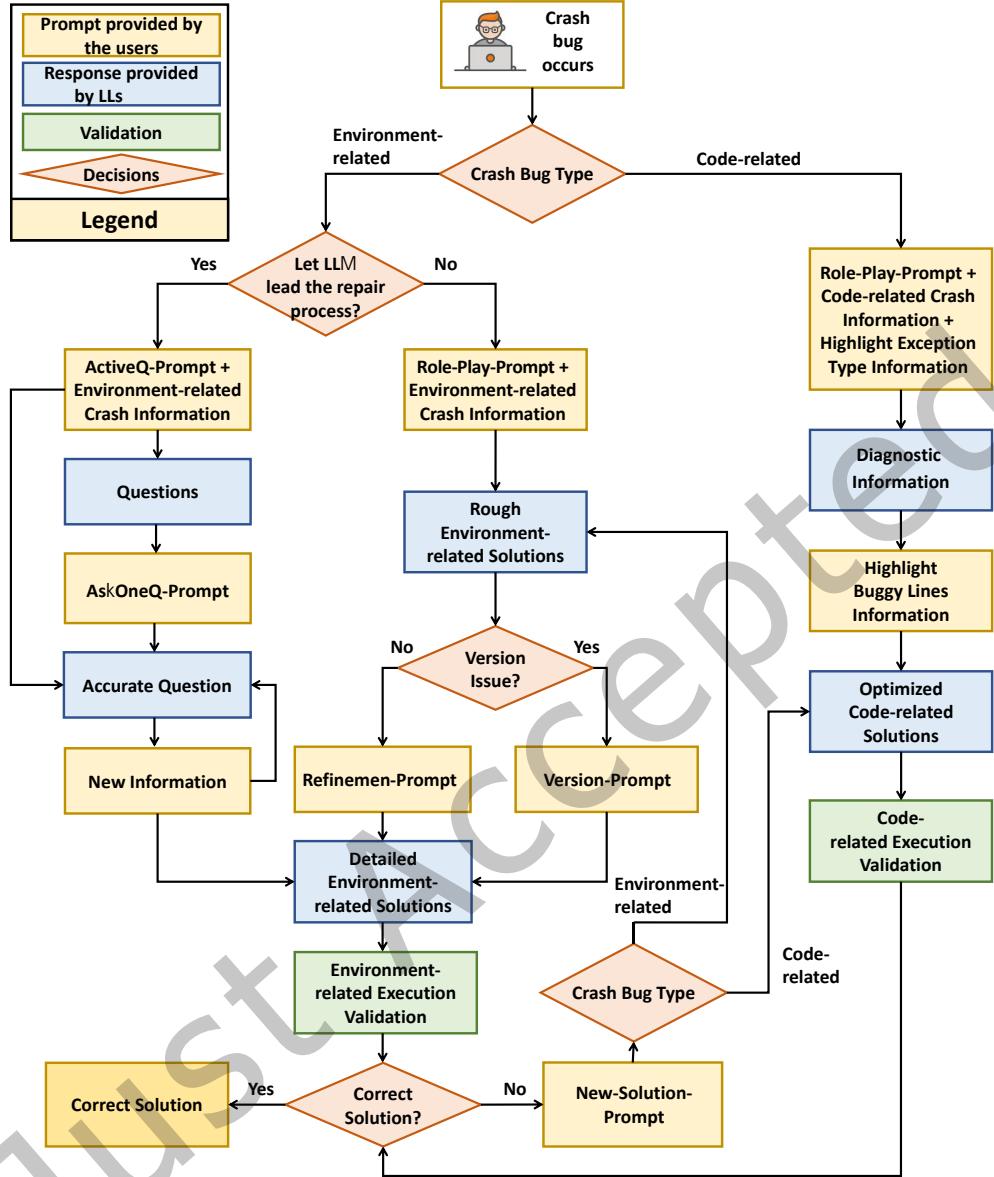


Fig. 3. Detailed Workflow of IntDiagSolver with Specific Prompts

Code-related Crash Bug Resolution. The main challenge for resolving code-related crash bugs lies in accurately diagnosis the root cause of crash bugs (refer to finding 3). Therefore, utilizing a Chain-of-Thought strategy to divide the resolution into diagnosis and repair stages is beneficial. This strategy entails deploying prompts to guide LLMs in identifying the buggy line, subsequently emphasizing the pinpointed information for efficient repair. Additionally, accentuating the exception type and providing a comprehensive crash context are instrumental in amplifying the effectiveness of the resolving (refer to finding 2).

Environment-related Crash Bug Resolution. For environment-related crash bugs resolving, users can choose to let LLMs guide the repair process. This is helpful for beginners with limited crash bug knowledge, while seasoned developers may not need LLMs' active guidance. Moreover, version issues present unique challenges as discussed in section 2.4 and finding 13. If users choose not to rely on LLMs' guidance, they should carefully examine the initial solution to identify any version issues. Version prompts should be utilized to elicit specific details in such cases.

Validation Process. Validation involves receiving LLMs' solutions and confirming whether they fix the crash bug. If not, new solution prompts are used, with a maximum retry limit (e.g., 3 attempts in our setting).

5 Evaluation of IntDiagSolver

To assess the effectiveness and generalizability of IntDiagSolver, we conducted comprehensive assessments for resolving crash bugs utilizing diverse LLMs. Initially, we conducted experiments by integrating IntDiagSolver based on GPT-3.5, which we employed in RQ1 and RQ2 (as detailed in Section 5.1). Since IntDiagSolver was developed based on the insights from RQ1 and RQ2, our objective was to determine its efficacy in addressing the identified challenges and previously unresolved crash bugs from those studies. Moreover, we further evaluated the generalizability of IntDiagSolver by applying it to a range of state-of-the-art LLMs, encompassing both open-source and closed-source models (Section 5.2), as well as by testing its effectiveness on crash bugs from different programming languages (Section 5.3).

5.1 RQ3: Effectiveness of IntDiagSolver

We assess the effectiveness of IntDiagSolver in resolving crash bugs using GPT-3.5 and compare it to a baseline scenario that employed with a basic prompt.

5.1.1 Participants. We invited two participants, each with over two years of Java programming experience, to conduct experiments and evaluations independently. Their evaluations were based on the criteria for the correctness of localization and repair as RQ1-defined metrics. In cases of disagreement, an author involved in RQ1 and RQ2 provided arbitration.

5.1.2 Benchmark. From the initial 100-crash-bug benchmark in Section 2.1, we deliberately chose unrepairs instances (see Table 2) despite receiving complete crash descriptions. This selection forms a refined benchmark, comprising 11 code-related crash bugs and 30 environment-related crash bugs crashes, totaling 41 cases.

5.1.3 Metrics. The experiment's metrics align with those presented in Section 2.3.2. Additionally, to assess the effectiveness of continuous interaction, we record the number of interaction rounds required to achieve the final correct localization or repair result. To mitigate any potential concerns with the randomness of the LLM, each query was executed three times using both IntDiagSolver and the baseline methods. For the calculation of localization and repair accuracy, a case is considered successfully resolved only if at least two out of the three runs yield a correct result.

5.1.4 Results. The results demonstrate significant improvements in resolving both code-related crash bugs and environment-related crash bugs. Specifically, localization accuracy increased by 36.3% and 43.3%, while repair accuracy improved by 36.4% and 53.3%, respectively.

In code-related crash bugs resolution, 4 additional crash bugs were effectively localized and repaired. Three among these were repaired by emphasizing the buggy line in subsequent interactions, while highlighting the exception type in initial crash description facilitated one's localization. Concerning environment-related crash bugs, the improvement was more significant: 13 more bugs were accurately localized, and 16 more were

Table 5. Effectiveness of IntDiagSolver across different LLMs

Model	Type	Prompt	Localization	Repair	Multi. Rounds	Localization Accuracy	Repair Accuracy
GPT-3.5	Code Related	Basic Prompt	3/11	0/11	/	27.3%	0.0%
		IntDiagSolver	7/11 (+4)	4/11 (+4)	3	63.6% (+36.3%)	36.4% (+36.4%)
	Env. Related	Basic Prompt	12/30	0/30	/	40.0%	0.0%
		IntDiagSolver	25/30 (+13)	16/30 (+16)	14	83.3% (+43.3%)	53.3% (+53.3%)
GPT-4	Code Related	Basic Prompt	8/11	5/11	/	72.7%	45.5%
		IntDiagSolver	9/11 (+1)	6/11 (+1)	1	81.8% (+9.1%)	54.5% (+9.1%)
	Env. Related	Basic Prompt	21/30	10/30	/	70.0%	33.3%
		IntDiagSolver	25/30 (+4)	21/30 (+11)	9	83.3% (+13.3%)	70.0% (+36.7%)
Claude	Code Related	Basic Prompt	5/11	4/11	/	45.5%	36.4%
		IntDiagSolver	6/11 (+1)	5/11 (+1)	1	54.5% (+9.1%)	45.5% (+9.1%)
	Env. Related	Basic Prompt	12/30	5/30	/	40.0%	16.7%
		IntDiagSolver	18/30 (+6)	13/30 (+8)	8	60.0% (+20.0%)	43.3% (+26.6%)
Codellama	Code Related	Basic Prompt	4/11	2/11	/	36.4%	18.2%
		IntDiagSolver	6/11 (+2)	4/11 (+2)	1	54.5% (+18.1%)	36.4% (+18.2%)
	Env. Related	Basic Prompt	7/30	2/30	/	23.3%	6.7%
		IntDiagSolver	12/30 (+5)	9/30 (+7)	7	40.0% (+16.7%)	30.0% (+23.3%)
Deepseek-r1	Code Related	Basic Prompt	8/11	7/11	/	72.7%	63.6%
		IntDiagSolver	11/11 (+3)	11/11 (+4)	2	100.0% (+27.3%)	100.0% (+36.4%)
	Env. Related	Basic Prompt	19/30	13/30	/	63.3%	43.3%
		IntDiagSolver	28/30 (+9)	26/30 (+13)	12	93.3% (+30.0%)	86.7% (+43.4%)
Qwen3-coder	Code Related	Basic Prompt	8/11	7/11	/	72.7%	63.6%
		IntDiagSolver	10/11 (+2)	9/11 (+2)	1	90.9% (+18.2%)	81.8% (+18.2%)
	Env. Related	Basic Prompt	15/30	11/30	/	50.0%	36.7%
		IntDiagSolver	26/30 (+11)	24/30 (+13)	10	86.7% (+36.7%)	80.0% (+43.3%)

successfully repaired. Furthermore, 14 crash bugs were resolved through iterative interactions, highlighting the efficacy of our IntDiagSolver framework.

5.2 RQ4.a: Generalizability of IntDiagSolver across LLMs

We evaluate IntDiagSolver’s generalizability on five additional LLMs, employing the same benchmarks and metrics as in RQ3 (Section 5.1).

5.2.1 Studied LLMs. We select five state-of-the-art LLMs: two closed-source models (GPT-4 [59] and Claude [5]) and three open-source models (Codellama-34b [64], DeepSeek-R1 [27] and Qwen3-Coder-Instruction [68]). To ensure experimental consistency, we maintain a uniform approach by conducting manual interactions through a web interface for all the studied LLMs.

5.2.2 Results. Overall Performance. Table 5 demonstrates IntDiagSolver’s remarkable generalizability, effectively improving resolution for both code-related crash bugs and environment-related crash bugs across various closed-source and open-source LLMs when compared to the basic prompt, with substantial enhancements ranging from 9.1% to 43.3% in localization and 9.1% to 53.3% in repair. **Localization vs. Repair.** Overall, the improvement in the repair stage is more pronounced than in the localization stage, particularly for environment-related crash bugs. This highlights the effectiveness of IntDiagSolver in addressing the challenge where LLMs often struggle to provide specific repair steps for environment-related crash bugs (see Finding 6). **Weaker LLMs vs. Stronger LLMs.** All models benefit significantly from the introduction of IntDiagSolver. For code-related crash bugs, some stronger models (such as DeepSeek-R1) already achieve high performance with the basic prompt, sometimes even outperforming weaker models (such as GPT-3.5 and CodeLlama) using IntDiagSolver. However, for environment-related crash bugs, weaker models using IntDiagSolver still surpass the performance of stronger models using only the basic prompt. This indicates that resolving environment-related crash bugs remains challenging even for state-of-the-art LLMs, underscoring the value of IntDiagSolver. Additionally, except for CodeLlama, all models

show greater improvement in resolving environment-related crash bugs than code-related crash bugs, which is consistent with our focus on environment-related bugs in Section 3 and leads to more substantial enhancements in this area.

5.3 RQ4.b: Generalizability of IntDiagSolver across Programming Languages

We constructed an expanded small-scale multilingual dataset to evaluate the generalizability of IntDiagSolver across different programming languages, including Java, C/C++, and Python, using the same metrics as in RQ3 (Section 5.1).

5.3.1 Benchmark. To build this multilingual dataset, we invited the same two participants from Section 5.2 to follow the data collection procedures described in Section 2.1. We collected a total of 42 environment-related crash bugs spanning Java, C/C++, and Python for further generalization evaluation (with 10 cases for Java, 19 for Python, and 6 for C/C++). To further assess the effectiveness of LLMs on previously unseen data, all 42 crash bugs were sourced from Stack Overflow posts published after 2024.

5.3.2 Studied LLMs. For this generalization experiment, we selected two representative advanced models: a closed-source model (i.e., GPT-4o [15]) and a open-source model(i.e., DeepSeek-R1 [27]).

5.3.3 Results. Overall Performance. As shown in Table 6, IntDiagSolver consistently demonstrates significant improvements over the baselines across all programming languages. Specifically, localization accuracy improved by 12.0% to 16.6%, and repair accuracy improved by 19.0% to 21.4% across the two advanced LLMs. **Impact of Programming Language.** Across all three programming languages, IntDiagSolver shows varying degrees of improvement when combined with both LLMs. Notably, both models achieved the most significant gains in resolving Python crash bugs, particularly in repair accuracy. **Performance on Unseen Data.** Overall, the comparative results between IntDiagSolver and the baseline on this unseen dataset closely mirror those observed in RQ4.a. The application of IntDiagSolver results in a significant performance improvement, which not only confirms its effectiveness on previously unseen data, but also demonstrates that the experimental results in RQ1–RQ4 are not influenced by data contamination.

Table 6. Effectiveness of IntDiagSolver across different programming languages

Model	Prompt	Language	Localization	Repair	Localization Accuracy	Repair Accuracy
GPT-4o	Basic Prompt	Python	15/20	11/20	75.0%	55.0%
		Java	8/13	7/13	61.5%	53.8%
		C/C++	5/9	4/9	55.6%	44.4%
		All	28/42	22/42	66.7%	52.4%
	IntDiagSolver	Python	19/20 (+4)	18/20 (+7)	95.0% (+20.0%)	90.0% (+35.0%)
		Java	10/13 (+2)	7/13 (+0)	76.9% (+15.4%)	53.8% (+0%)
		C/C++	6/9 (+1)	5/9 (+1)	66.7% (+11.1%)	55.6% (+11.2%)
		All	35/42 (+7)	30/42 (+8)	83.3% (+16.6%)	71.4% (+19.0%)
DeepSeek-r1	Basic Prompt	Python	16/20	11/20	80.0%	55.0%
		Java	8/13	6/13	61.5%	46.2%
		C/C++	5/9	5/9	55.6%	55.6%
		All	29/42	22/42	69.0%	52.4%
	IntDiagSolver	Python	18/20 (+2)	18/20 (+7)	90.0% (+10.0%)	90.0% (+35.0%)
		Java	10/13 (+2)	8/13 (+2)	76.9% (+15.4%)	61.5% (+15.3%)
		C/C++	6/9 (+1)	5/9 (+0)	66.7% (+11.1%)	55.6% (+0%)
		All	34/42 (+5)	31/42 (+9)	81.0% (+12.0%)	73.8% (+21.4%)

6 Related Work

6.1 LLMs for Software Engineering

In recent years, LLMs have gained considerable attention across various research domains, including mathematics [26, 34], education [18, 36], and natural language processing [19, 40]. In the field of software engineering (SE), LLMs have demonstrated their potential by being applied to a wide range of tasks [50], from code generation [30, 51, 60, 86] and code summarization [17, 72] to software maintenance tasks, including vulnerability detection [67, 80, 87], test generation [81, 84], code translation [82, 85], and program repair [25, 38, 41, 56, 62, 66, 77]. This broad SE application stems from their robust training on extensive code and text data, enhancing both linguistic and code comprehension.

6.2 LLMs for Resolving Crash Bugs

Software crashes pose an enduring challenge in software development, driving research in areas such as crash reproduction [42, 58], crash localization [32, 63, 73, 76], and crash repair [35, 52, 71]. Efforts to tackle code-related crash bugs have resulted in research on automatic localization [23, 24, 78] and repair [22, 56, 77]. In contrast, environment-related crash bugs, due to their diverse origins, often rely on solutions from online Q&A forums like SO [31, 52, 54, 71].

Recent surveys on LLMs in software engineering [37, 83] have explored their applications, performance, and challenges, including bug localization and program repair. Preliminary studies [41, 66, 77] on ChatGPT’s potential in addressing code-related bugs have primarily focused on simple scenarios with fixed prompts. In contrast, our study: (i) diverges from existing work by categorizing the resolution process into two stages: localization and repair, (ii) comprehensively assesses ChatGPT’s effectiveness in resolving real-world crash bugs sourced from SO, (iii) explores the effects of different prompts in various interaction strategies, and (iv) notably extends beyond code-related problems to include crash bugs caused by external environmental factors, an aspect that has received limited attention in prior research.

6.3 Prompt Engineering for Software Engineering

Prompt engineering is an emerging discipline that optimizes prompts for various applications of LLMs across different domains. Methodologies in this field include Few-Shot Prompting [21], Chain-of-Thought Prompting (CoT) [74, 88], Tree-of-Thoughts (ToT) [79], and Knowledge Prompting [49]. In SE domain, prompt engineering enhances LLMs in various tasks. Researchers have used CoT Prompting to improve code generation [45, 48], explored argumentation and task decomposition strategies for better unit test generation [84], and combined CoT Prompting with static analysis for more efficient vulnerability detection [87]. In this study, we explored various prompt templates and advanced techniques to effectively resolve crash bugs. We further propose IntDiagSolver, employing Knowledge Prompting to resolve crash bugs. It breaks the process into localization and repair with multi-turn interactions involving LLMs, following the principles of CoT Prompting.

7 Threats of Validity

The empirical study used a dataset of 100 Java-related crash bugs from SO, which might limit generalizability. To address this, we constructed an expanded small-scale multilingual dataset to evaluate the generalizability of our method across different programming languages in RQ4. Another threat arises from the fact that our evaluation compares LLM-generated answers against the accepted answers on Stack Overflow, without additional empirical validation. Nevertheless, the accepted answers were originally verified for correctness by the question askers themselves, providing a reasonable basis for reliability. Moreover, due to the lack of sufficient contextual information in most Stack Overflow threads, reproducing the crashes for empirical validation was infeasible, so we followed the same evaluation methodology adopted in prior studies [53, 55, 70?]. For future work, we plan to

manually collect and construct a dedicated benchmark that includes sufficient contextual information to enable full empirical validation of crash reproduction and resolution.

In our exploratory study (RQ1 and RQ2), employing a specific LLM (i.e., GPT-3.5) may limit the generalizability of our findings and the broader applicability of the proposed IntDiagSolver approach. To mitigate this concern, in RQ3, we evaluated our approach using various state-of-the-art LLMs, reinforcing our initial findings and indicating wider applicability across different models. Another potential limitation of RQ2 stems from exploring interactions using only 10 cases, which may introduce selection bias. However, subsequent practical evaluations demonstrated promising effectiveness across a broader range of cases, as shown in Table 5. Due to the inherently stochastic nature of LLM outputs, there is an inherent risk to the reliability and validity of experimental results. To address this, we conducted multiple trials and rigorously documented all interactions with LLMs to ensure transparency and reproducibility [16].

An additional threat involves potential data leakage from the Stack Overflow benchmark dataset. Nevertheless, our investigations in RQ1 and RQ2 revealed that model performance predominantly depends on contextual variations of crash bug scenarios, as well as the diverse prompt formulations and interaction strategies used. This indicates that the model is not simply memorizing the data used in our study, even if the LLM had been exposed to this data. Furthermore, the effectiveness on previously unseen dataset demonstrates the experimental results in RQ1–RQ4 are not influenced by data contamination.

8 Conclusion

This study first empirically investigates the effectiveness of LLMs in resolving real-world environment-related crash bugs by comparing their performance with code-related crash bugs, and exploring the impact of different prompt strategies and multi-round interactions. Leveraging experimental findings, we introduce IntDiagSolver, a methodology that optimizes the interaction process and prompt design for resolving crash bugs. Our evaluation results demonstrate that IntDiagSolver achieves significant improvements in resolving environment-related crash bugs across multiple LLMs through its optimized interaction process.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62332005).

References

- [1] 2021. *Stack Overflow data dump version from September 4, 2021*. Retrieved May 4, 2022 from <https://archive.org/download/stackexchange/>
- [2] 2022. *Libaries.io open data*. Retrieved January 20, 2023 from <https://libraries.io/data>
- [3] 2022. *Maven Central Repository*. Retrieved January 20, 2023 from <https://mvnrepository.com>
- [4] 2023. *chatgpt-prompts*. <https://github.com/f/awesome-chatgpt-prompts>
- [5] 2023. *Claude*. <https://www.anthropic.com/index/introducing-claude>
- [6] 2023. *JDK 1.8*. [https://docs.oracle.com/javase/8/docs/api/overview-summary.html/](https://docs.oracle.com/javase/8/docs/api/overview-summary.html)
- [7] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/68199510>
- [8] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/65084069>
- [9] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/22928450>
- [10] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/28097042>
- [11] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/24305296>
- [12] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/49871007>
- [13] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/43320334>
- [14] 2023. *the Stack Overflow post ID of the crash bug*. <https://stackoverflow.com/questions/37771758>
- [15] 2024. *GPT-4-turbo*. <https://platform.openai.com/docs/models/gpt-4-and-gpt-4-turbo>
- [16] 2024. *replication package*. <https://llm4crasolver.github.io>
- [17] Toufique Ahmed and Premkumar Devanbu. 2023. Few-Shot Training LLMs for Project-Specific Code-Summarization. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering* (Rochester, MI, USA) (ASE '22). Association for Computing Machinery, New York, NY, USA, Article 177, 5 pages. <https://doi.org/10.1145/3551349.3559555>

- [18] David Baidoo-Anu and Leticia Owusu Ansah. 2023. Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Available at SSRN 4337484* (2023).
- [19] Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holly Lovenia, Ziwei Ji, Tiezheng Yu, Willy Chung, Quyet V. Do, Yan Xu, and Pascale Fung. 2023. A Multitask, Multilingual, Multimodal Evaluation of ChatGPT on Reasoning, Hallucination, and Interactivity. *arXiv:2302.04023 [cs.CL]*
- [20] Islem Bouzenia, Premkumar Devanbu, and Michael Pradel. 2024. Repairagent: An autonomous, llm-based agent for program repair. *arXiv preprint arXiv:2403.17134* (2024).
- [21] 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.
- [22] Jialun Cao, Meiziniu Li, Ming Wen, and Shing-chi Cheung. 2023. A study on Prompt Design, Advantages and Limitations of ChatGPT for Deep Learning Program Repair. *arXiv preprint arXiv:2304.08191* (2023).
- [23] Junming Cao, Shouliang Yang, Wenhui Jiang, Hushuang Zeng, Beijun Shen, and Hao Zhong. 2020. Bugpecker: Locating faulty methods with deep learning on revision graphs. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*. 1214–1218.
- [24] Agnieszka Ciborowska and Kostadin Damevski. 2022. Fast Changeset-Based Bug Localization with BERT. In *Proceedings of the 44th International Conference on Software Engineering* (Pittsburgh, Pennsylvania) (ICSE '22). Association for Computing Machinery, New York, NY, USA, 946–957. <https://doi.org/10.1145/3510003.3510042>
- [25] Dekun Dai, MingWei Liu, Anji Li, Jialun Cao, Yanlin Wang, Chong Wang, Xin Peng, and Zibin Zheng. 2025. FeedbackEval: A Benchmark for Evaluating Large Language Models in Feedback-Driven Code Repair Tasks. *arXiv:2504.06939 [cs.SE]* <https://arxiv.org/abs/2504.06939>
- [26] Ernest Davis. 2023. Mathematics, word problems, common sense, and artificial intelligence. *arXiv:2301.09723 [cs.AI]*
- [27] DeepSeek-AI. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv:2501.12948 [cs.CL]* <https://arxiv.org/abs/2501.12948>
- [28] Tejinder Dhaliwal, Foutse Khomh, and Ying Zou. 2011. Classifying field crash reports for fixing bugs: A case study of Mozilla Firefox. In *IEEE 27th International Conference on Software Maintenance, ICSM 2011, Williamsburg, VA, USA, September 25-30, 2011*. IEEE Computer Society, 333–342. <https://doi.org/10.1109/ICSM.2011.6080800>
- [29] Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. 2023. Self-collaboration Code Generation via ChatGPT. *arXiv preprint arXiv:2304.07590* (2023).
- [30] Xueying Du, Mingwei Liu, Kaixin Wang, Hanlin Wang, Junwei Liu, Yixuan Chen, Jiayi Feng, Chaofeng Sha, Xin Peng, and Yiling Lou. 2024. Evaluating large language models in class-level code generation. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*. 1–13.
- [31] Xueying Du, Yiling Lou, Mingwei Liu, Xin Peng, and Tianyong Yang. 2023. KG4CraSolver: Recommending Crash Solutions via Knowledge Graph. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1242–1254.
- [32] Xueying Du, Geng Zheng, Kaixin Wang, Jiayi Feng, Wentai Deng, Mingwei Liu, Bihuan Chen, Xin Peng, Tao Ma, and Yiling Lou. 2024. Vul-RAG: Enhancing LLM-based Vulnerability Detection via Knowledge-level RAG. *arXiv preprint arXiv:2406.11147* (2024).
- [33] Rudolf Ferenc, Zoltán Tóth, Gergely Ladányi, István Siket, and Tibor Gyimóthy. 2018. A public unified bug dataset for java. In *Proceedings of the 14th international conference on predictive models and data analytics in software engineering*. 12–21.
- [34] Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Christian Petersen, Alexis Chevalier, and Julius Berner. 2023. Mathematical Capabilities of ChatGPT. *arXiv:2301.13867 [cs.LG]*
- [35] Qing Gao, Hansheng Zhang, Jie Wang, Yingfei Xiong, Lu Zhang, and Hong Mei. 2015. Fixing Recurring Crash Bugs via Analyzing Q&A Sites (T). In *30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA, November 9-13, 2015*. Myra B. Cohen, Lars Grunske, and Michael Whalen (Eds.). IEEE Computer Society, 307–318. <https://doi.org/10.1109/ASE.2015.81>
- [36] Aidan Gilson, Conrad Safranek, Thomas Huang, Vimig Socrates, Ling Chi, Richard Andrew Taylor, and David Chartash. 2022. How Well Does ChatGPT Do When Taking the Medical Licensing Exams? The Implications of Large Language Models for Medical Education and Knowledge Assessment. *medRxiv* (2022), 2022–12.
- [37] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2023. Large language models for software engineering: A systematic literature review. *arXiv preprint arXiv:2308.10620* (2023).
- [38] Nan Jiang, Kevin Liu, Thibaud Lutellier, and Lin Tan. 2023. Impact of Code Language Models on Automated Program Repair. *CoRR* abs/2302.05020 (2023). <https://doi.org/10.48550/arXiv.2302.05020> *arXiv:2302.05020*
- [39] Xue Jiang, Yihong Dong, Lecheng Wang, Qiwei Shang, and Ge Li. 2023. Self-planning code generation with large language model. *arXiv preprint arXiv:2303.06689* (2023).
- [40] Wenxiang Jiao, Wenxuan Wang, Jen-tse Huang, Xing Wang, and Zhaopeng Tu. 2023. Is ChatGPT a good translator? A preliminary study. *arXiv preprint arXiv:2301.08745* (2023).

[41] Matthew Jin, Syed Shahriar, Michele Tufano, Xin Shi, Shuai Lu, Neel Sundaresan, and Alexey Svyatkovskiy. 2023. Inferfix: End-to-end program repair with llms. *arXiv preprint arXiv:2303.07263* (2023).

[42] Wei Jin and Alessandro Orso. 2012. BugRedux: Reproducing field failures for in-house debugging. In *34th International Conference on Software Engineering, ICSE 2012, June 2-9, 2012, Zurich, Switzerland*, Martin Glinz, Gail C. Murphy, and Mauro Pezzè (Eds.). IEEE Computer Society, 474–484. <https://doi.org/10.1109/ICSE.2012.6227168>

[43] René Just, Darioush Jalali, and Michael D. Ernst. 2014. Defects4J: A Database of existing faults to enable controlled testing studies for Java programs. In *ISSTA 2014, Proceedings of the 2014 International Symposium on Software Testing and Analysis*. San Jose, CA, USA, 437–440. Tool demo.

[44] Rafael-Michael Karampatsis and Charles Sutton. 2020. How often do single-statement bugs occur? the manyssubsts4j dataset. In *Proceedings of the 17th International Conference on Mining Software Repositories*. 573–577.

[45] Jia Li, Ge Li, Yongmin Li, and Zhi Jin. 2023. Structured Chain-of-Thought Prompting for Code Generation. *arXiv preprint arXiv:2305.06599* (2023).

[46] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).

[47] 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* (Vancouver, BC, Canada) (*SPLASH Companion 2017*). Association for Computing Machinery, New York, NY, USA, 55–56. <https://doi.org/10.1145/3135932.3135941>

[48] Chao Liu, Xuanlin Bao, Hongyu Zhang, Neng Zhang, Haibo Hu, Xiaohong Zhang, and Meng Yan. 2023. Improving ChatGPT Prompt for Code Generation. *arXiv preprint arXiv:2305.08360* (2023).

[49] Jiacheng Liu, Alisa Liu, Ximing Lu, Sean Welleck, Peter West, Ronan Le Bras, Yejin Choi, and Hannaneh Hajishirzi. 2022. Generated knowledge prompting for commonsense reasoning. *arXiv preprint arXiv:2110.08387* (2022).

[50] Junwei Liu, Kaixin Wang, Yixuan Chen, Xin Peng, Zhenpeng Chen, Lingming Zhang, and Yiling Lou. 2024. Large Language Model-Based Agents for Software Engineering: A Survey. *CoRR* abs/2409.02977 (2024). <https://doi.org/10.48550/ARXIV.2409.02977> arXiv:2409.02977

[51] Junwei Liu, Chen Xu, Chong Wang, Tong Bai, Weitong Chen, Kaseng Wong, Yiling Lou, and Xin Peng. 2025. EvoDev: An Iterative Feature-Driven Framework for End-to-End Software Development with LLM-based Agents. *CoRR* abs/2511.02399 (2025). <https://doi.org/10.48550/ARXIV.2511.02399> arXiv:2511.02399

[52] Sonal Mahajan, Negarsadat Abolhassani, and Mukul R. Prasad. 2020. Recommending stack overflow posts for fixing runtime exceptions using failure scenario matching. In *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*. ACM, 1052–1064.

[53] Sonal Mahajan, Negarsadat Abolhassani, and Mukul R. Prasad. 2020. Recommending stack overflow posts for fixing runtime exceptions using failure scenario matching. In *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*. ACM, 1052–1064. <https://doi.org/10.1145/3368089.3409764>

[54] Sonal Mahajan and Mukul R. Prasad. 2022. Providing Real-time Assistance for Repairing Runtime Exceptions using Stack Overflow Posts. In *15th IEEE Conference on Software Testing, Verification and Validation, ICST 2022, Valencia, Spain, April 4-14, 2022*. IEEE, 196–207.

[55] Sonal Mahajan and Mukul R. Prasad. 2022. Providing Real-time Assistance for Repairing Runtime Exceptions using Stack Overflow Posts. In *15th IEEE Conference on Software Testing, Verification and Validation, ICST 2022, Valencia, Spain, April 4-14, 2022*. IEEE, 196–207. <https://doi.org/10.1109/ICST53961.2022.00030>

[56] Ehsan Mashhadi and Hadi Hemmati. 2021. Applying CodeBERT for Automated Program Repair of Java Simple Bugs. In *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. 505–509. <https://doi.org/10.1109/MSR52588.2021.00063>

[57] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia Medica: Biochemia Medica* 22, 3 (2012), 276–282.

[58] Mathieu Nayrolles, Abdelwahab Hamou-Lhadj, Sofiène Tahar, and Alf Larsson. 2017. A bug reproduction approach based on directed model checking and crash traces. *J. Softw. Evol. Process.* 29, 3 (2017). <https://doi.org/10.1002/smri.1789>

[59] OpenAI. 2023. GPT-4 Technical Report. *CoRR* abs/2303.08774 (2023). <https://doi.org/10.48550/arXiv.2303.08774> arXiv:2303.08774

[60] Gabriel Poesia, Oleksandr Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani. 2022. Synchromesh: Reliable code generation from pre-trained language models. *arXiv:2201.11227* [cs.LG]

[61] Julian Aron Prenner and Romain Robbes. 2021. Automatic Program Repair with OpenAI’s Codex: Evaluating QuixBugs. *arXiv:2111.03922* [cs.SE]

[62] Alfin Wijaya Rahardja, Junwei Liu, Weitong Chen, Zhenpeng Chen, and Yiling Lou. 2025. Can Agents Fix Agent Issues? *CoRR* abs/2505.20749 (2025). <https://doi.org/10.48550/ARXIV.2505.20749> arXiv:2505.20749

[63] Mohammad Masudur Rahman and Chanchal K. Roy. 2018. Improving IR-based bug localization with context-aware query reformulation. In *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*, Gary T. Leavens, Alessandro Garcia, and Corina S. Pasareanu (Eds.). ACM, 621–632. <https://doi.org/10.1145/3236024.3236065>

[64] Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémie Rapin, et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950* (2023).

- [65] Dominik Sobania, Martin Briesch, Carol Hanna, and Justyna Petke. 2023. An analysis of the automatic bug fixing performance of chatgpt. *arXiv preprint arXiv:2301.08653* (2023).
- [66] Dominik Sobania, Martin Briesch, Carol Hanna, and Justyna Petke. 2023. An Analysis of the Automatic Bug Fixing Performance of ChatGPT. *arXiv:2301.08653 [cs.SE]*
- [67] Yuqiang Sun, Daoyuan Wu, Yue Xue, Han Liu, Haijun Wang, Zhengzi Xu, Xiaofei Xie, and Yang Liu. 2023. When GPT Meets Program Analysis: Towards Intelligent Detection of Smart Contract Logic Vulnerabilities in GPTScan. *arXiv:2308.03314 [cs.CR]*
- [68] Qwen Team. 2025. Qwen3 Technical Report. *arXiv:2505.09388 [cs.CL]* <https://arxiv.org/abs/2505.09388>
- [69] Huiyan Wang, Shuguan Liu, Lingyu Zhang, and Chang Xu. 2023. Automatically Resolving Dependency-Conflict Building Failures via Behavior-Consistent Loosening of Library Version Constraints. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 198–210.
- [70] Haoye Wang, Xin Xia, David Lo, John Grundy, and Xinyu Wang. 2021. Automatic solution summarization for crash bugs. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. IEEE, 1286–1297.
- [71] Haoye Wang, Xin Xia, David Lo, John C. Grundy, and Xinyu Wang. 2021. Automatic Solution Summarization for Crash Bugs. In *43rd IEEE/ACM International Conference on Software Engineering, ICSE 2021, Madrid, Spain, 22-30 May 2021*. IEEE, 1286–1297.
- [72] Ruyun Wang, Hanwen Zhang, Guoliang Lu, Lei Lyu, and Chen Lyu. 2020. Fret: Functional reinforced transformer with bert for code summarization. *IEEE Access* 8 (2020), 135591–135604.
- [73] Shaohua Wang, Foutse Khomh, and Ying Zou. 2013. Improving bug localization using correlations in crash reports. In *Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13, San Francisco, CA, USA, May 18–19, 2013*, Thomas Zimmermann, Massimiliano Di Penta, and Sunghun Kim (Eds.). IEEE Computer Society, 247–256. <https://doi.org/10.1109/MSR.2013.6624036>
- [74] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* 35 (2022), 24824–24837.
- [75] Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382* (2023).
- [76] Rongxin Wu, Hongyu Zhang, Shing-Chi Cheung, and Sunghun Kim. 2014. CrashLocator: locating crashing faults based on crash stacks. In *International Symposium on Software Testing and Analysis, ISSTA '14, San Jose, CA, USA - July 21 - 26, 2014*, Corina S. Pasareanu and Darko Marinov (Eds.). ACM, 204–214. <https://doi.org/10.1145/2610384.2610386>
- [77] Chunqiu Steven Xia and Lingming Zhang. 2023. Conversational automated program repair. *arXiv preprint arXiv:2301.13246* (2023).
- [78] Zhou Yang, Jieke Shi, Shaowei Wang, and David Lo. 2021. Inclb: Incremental bug localization. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 1223–1226.
- [79] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601* (2023).
- [80] Alperen Yıldız, Sin G. Teo, Yiling Lou, Yebo Feng, Chong Wang, and Dinal Mon Divakaran. 2025. Benchmarking LLMs and LLM-based Agents in Practical Vulnerability Detection for Code Repositories. *CoRR* abs/2503.03586 (2025). <https://doi.org/10.48550/ARXIV.2503.03586> arXiv:2503.03586
- [81] Shengcheng Yu, Chunrong Fang, Yuchen Ling, Chentian Wu, and Zhenyu Chen. 2023. LLM for Test Script Generation and Migration: Challenges, Capabilities, and Opportunities. *arXiv:2309.13574 [cs.SE]*
- [82] Zhiqiang Yuan, Weitong Chen, Hanlin Wang, Kai Yu, Xin Peng, and Yiling Lou. 2024. TRANSAGENT: An LLM-Based Multi-Agent System for Code Translation. *CoRR* abs/2409.19894 (2024). <https://doi.org/10.48550/ARXIV.2409.19894> arXiv:2409.19894
- [83] Zhiqiang Yuan, Junwei Liu, Qiancheng Zi, Mingwei Liu, Xin Peng, and Yiling Lou. 2023. Evaluating instruction-tuned large language models on code comprehension and generation. *arXiv preprint arXiv:2308.01240* (2023).
- [84] Zhiqiang Yuan, Yiling Lou, Mingwei Liu, Shiji Ding, Kaixin Wang, Yixuan Chen, and Xin Peng. 2023. No More Manual Tests? Evaluating and Improving ChatGPT for Unit Test Generation. *arXiv:2305.04207 [cs.SE]*
- [85] Zhiqiang Yuan, Wenjun Mao, Zhuo Chen, Xiyue Shang, Chong Wang, Yiling Lou, and Xin Peng. 2025. Project-Level C-to-Rust Translation via Synergistic Integration of Knowledge Graphs and Large Language Models. *CoRR* abs/2510.10956 (2025). <https://doi.org/10.48550/ARXIV.2510.10956> arXiv:2510.10956
- [86] Zhengran Zeng, Hanzhuo Tan, Haotian Zhang, Jing Li, Yuqun Zhang, and Lingming Zhang. 2022. An extensive study on pre-trained models for program understanding and generation. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*. 39–51.
- [87] Chenyuan Zhang, Hao Liu, Jiutian Zeng, Kejing Yang, Yuhong Li, and Hui Li. 2023. Prompt-Enhanced Software Vulnerability Detection Using ChatGPT. *arXiv:2308.12697 [cs.SE]*
- [88] Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493* (2022).
- [89] Zhongkai Zhao, Bonan Kou, Mohamed Yilmaz Ibrahim, Muham Chen, and Tianyi Zhang. 2023. Knowledge-Based Version Incompatibility Detection for Deep Learning. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (, San Francisco, CA, USA,) (ESEC/FSE 2023). Association for Computing Machinery, New York, NY, 1–11.

Just Accepted