Trae Agent: An LLM-based Agent for Software Engineering with Test-time Scaling

Trae Research
Beijing, China
opensource@mail.trae.ai

Abstract

Software issue resolution is a critical challenge in software engineering and has garnered increasing attention in recent years. With the rapid advancement of large language models (LLMs), substantial progress has been made in addressing real-world software engineering tasks. Recent studies have introduced ensemble reasoning techniques to enhance the performance of LLM-based issue resolution. However, existing prompting-based methods still face limitations in effectively exploring large ensemble spaces and lack the capacity for repository-level understanding, both of which constrain their overall effectiveness. In this paper, we propose `Trae Agent`, the first agent-based ensemble reasoning approach for repository-level issue resolution. `Trae Agent` formulates our goal as an optimal solution search problem and addresses two key challenges, i.e., large ensemble spaces and repository-level understanding, through modular agents for generation, pruning, and selection. We conduct extensive experiments using three leading LLMs on the widely-adopted SWE-bench benchmark, comparing `Trae Agent` against four state-of-the-art ensemble reasoning techniques. Experimental results demonstrate that `Trae Agent` consistently achieves superior performance, with an average improvement of 10.22% over all baselines in terms of Pass@1. `Trae Agent` *has achieved first place on the SWE-bench Verified leaderboard, with a notable Pass@1 score of 75.20%.* We are pleased to release `Trae Agent` as an open-source project to support the research community, with all resources available at https://github.com/bytedance/trae-agent.

1 Introduction

Software issue resolution refers to the automated handling of newly reported bugs or feature requests during software development and maintenance, aiming to ensure correct and reliable system behavior [Fakhoury et al., 2023, Jimenez et al., 2023, Guo et al., 2025]. In recent years, this task has garnered growing interest from both academia and industry, driven by its potential to reduce developer burden and significantly enhance productivity [Yang et al., 2024, Zhang et al., 2024b]. Large language models (LLMs) have shown remarkable capabilities in function-level code tasks such as code generation [Lin et al., 2025, Tian et al., 2025] and automated program repair [Fan et al., 2023, Bouzenia et al., 2025]. Despite this progress, LLMs continue to face major challenges in resolving complex repository-level software issues. For example, GPT-4o achieves a resolution rate of 92.70% on the function-level HumanEval benchmark [Liu et al., 2023], but only 11.99% on the repository-level SWE-bench benchmark [Yang et al., 2024]. This performance gap underscores the difficulty of real-world software issue resolution, which often requires global understanding of large codebases, cross-file reasoning, and the detection of subtle, multi-component bugs [Ruan et al., 2025, Xia et al., 2025]. These challenges limit the practical adoption of LLM-based techniques in real-world software engineering and pose a risk to software quality if unresolved. Bridging this gap

---

The list of authors is provided at the end of the technical report.

remains a critical research goal toward realizing robust, scalable, and reliable automated software issue resolution.

Recently, many techniques have been proposed to enhance the performance of LLMs in software issue resolution [Yang et al., 2024, Wang et al., 2025b, Zhang et al., 2024b]. These methods primarily focus on improving patch generation, by carefully designing agent architectures and integrating external tools to assist LLMs in reasoning and producing correct fixes. For example, OpenDevin[Wang et al., 2025b] incorporates a planning module based on issue descriptions and coordinates multiple tools within an agent system to guide patch generation. While these approaches improve individual patch generation, recent studies [Zhang et al., 2024a, Li et al., 2025, Chen and Flaherty, 2025] have revealed a key property of LLMs: although the overall resolution rate remains stable across repeated runs, the set of issues successfully resolved varies across executions. This variability arises from the large action space and complex reasoning trajectories required for repository-level issue resolution, which often lead LLMs to pursue divergent solution paths. This observation has led to the emergence of a new direction, i.e., **ensemble reasoning**, which builds on patch generation approaches by producing multiple independent candidate patches and selecting a consensus solution that is more robust than any individual output.

Existing ensemble reasoning methods are prompting-based, relying on LLMs not only to generate candidate patches, but also to perform selection among them. For example, Augment[Chen and Flaherty, 2025] uses the LLM-as-a-judge [Wang et al., 2025a] paradigm to prompt the model to compare each candidate against the issue description and choose the best match. DeiBase[Zhang et al., 2024a] prompts the LLM to provide justification and confidence scores for each candidate patch, and selects the highest-scoring one. While prompting-based ensemble reasoning techniques have shown promise, they face fundamental limitations when applied to complex repository-level issue resolution. These limitations become especially evident in the face of two key challenges: First, prompting-based methods *struggle with the optimal solution search in large ensemble spaces*. As the number of candidate patches grows, subtle behavioral differences between them become increasingly difficult to distinguish using LLMs alone. Since patch selection is typically performed through a single prompt, surface-level syntactic comparisons often fail to capture deeper semantic nuances. Second, these methods *lack the capacity for repository-level understanding*, which is essential for accurately selecting patches in real-world software systems. Issues and fixes frequently span multiple files and modules, requiring cross-file reasoning, contextual awareness, and verification of patch correctness within the broader codebase. Prompting-based approaches operate in a stateless, single-turn manner and lack persistent memory or tool integration, making it difficult to track dependencies, execute validation steps, or build a coherent global view of the repository.

To overcome the limitations of prompting-based ensemble reasoning approaches, we propose the first agent-based ensemble reasoning approach for repository-level issue resolution, called `Trae Agent`. This approach is designed to enhance LLM-based issue resolution by formulating it as an optimal solution search problem [Harman and Jones, 2001, Clarke et al., 2003] and addressing the two core challenges through a modular agent-based architecture. It consists of three key components: *patch generation*, *patch pruning*, and *patch selection*. To improve the ensemble diversity, `Trae Agent` leverages a novel coder agent that first generates diverse candidate patches in parallel. To mitigate the difficulty of optimal solution search, `Trae Agent` applies a hierarchical pruning strategy that combines patch deduplication and regression testing strategies to eliminate redundant or faulty patches, thereby reducing the ensemble space while preserving promising candidates. To address the need for repository-level understanding, `Trae Agent` employs a selector agent that simulates a real-world program comprehension process [Eisenbarth et al., 2001, Di Lucca and Di Penta, 2005], encompassing both static review and dynamic verification. Specifically, this agent iteratively gathers and analyzes relevant code snippets (such as those referenced in the issue description, modified by patches, or related through dependencies) to build static understanding, and collects execution traces from automatically generated tests to build dynamic understanding. Based on these insights, the agent applies a majority voting strategy to select the most plausible patch. Due to its generalizable and modular design, `Trae Agent` potentially provides a solid foundation for advancing ensemble reasoning in other complex software engineering tasks.

We conduct extensive experiments to evaluate the performance of `Trae Agent` using three state-of-the-art LLMs (i.e., Gemini 2.5 Pro [DeepMind, 2025], Claude 3.7 Sonnet [Anthropic, 2025b], and GPT-4.1 [OpenAI, 2025a]) on the widely-used software issue resolution benchmark (i.e., SWE-bench [Jimenez et al., 2023]). Our results show that `Trae Agent` consistently and significantly

outperforms four state-of-the-art ensemble reasoning baselines (i.e., Augment, Augment w/ Pruning, DeiBase, and DeiBase w/ Pruning) across all evaluation settings, demonstrating its capability to enhance LLM-based software issue resolution. In particular, `Trae Agent` achieves a Pass@1 improvement ranging from 5.83% to 14.60%, where Pass@1 denotes the proportion of generated patches that pass all golden tests. We further investigate the impact of ensemble size, a critical hyper-parameter in ensemble reasoning. The results reveal that `Trae Agent` not only outperforms all baselines at each ensemble size but also continues to improve with larger ensemble sizes, while baseline methods exhibit varying degrees of performance degradation. Additionally, we conduct comprehensive ablation studies by constructing five variants of `Trae Agent`, each omitting a major component (i.e., patch pruning, patch deduplication, regression testing, selector agent, and majority voting). The results confirm the effectiveness and necessity of each part of our framework.

The main contributions of this paper are as follows:

- **Novel Approach**: We propose `Trae Agent`, the first agent-based ensemble reasoning framework for repository-level issue resolution. `Trae Agent` formulates our goal as an optimal solution search problem and addresses two key challenges, i.e., large ensemble spaces and repository-level understanding, through modular agents for generation, pruning, and selection.

- **Extensive Evaluation**: We conduct extensive experiments on the SWE-bench benchmark using three state-of-the-art LLMs, comparing `Trae Agent` with four leading baselines. Results show that `Trae Agent` consistently improves Pass@1 by 5.83%∼14.60%, demonstrating its robust effectiveness across settings. Particularly, `Trae Agent` has achieved first place on the SWE-bench Verified leaderboard, with a notable Pass@1 score of 75.20%.

- **Public Artifact**: To support the research community and facilitate future research, we release all resources available at `https://github.com/bytedance/trae-agent`. Our open-source GitHub repository has attracted over 8,000 stars (as of July 2025), reflecting strong community interest and adoption.

2 Motivation

To motivate the potential of ensemble reasoning approaches and the necessity of the two key components in `Trae Agent`, i.e., patch pruning (Section 3.3) and patch selection (Section 3.4), we present a comprehensive empirical analysis.

Figure 1 shows the performance of two state-of-the-art ensemble reasoning techniques (Augment [Code, 2025] and DeiBase [Zhang et al., 2024a]) on the widely-used SWE-bench Verified [Jimenez et al., 2023] benchmark using three leading LLMs (Gemini 2.5 Pro [DeepMind, 2025], Claude 3.7 Sonnet [Anthropic, 2025b], and GPT-4.1 [OpenAI, 2025a]). In particular, the Mixture setting refers to a round-robin combination of these three LLMs for generating candidate patches, aiming to enhance the diversity of the ensemble space. Taking Figure 1(1) as an example, the x-axis denotes the ensemble size N (i.e., the number of candidate patches generated by Gemini 2.5 Pro), while the y-axis shows the Pass@1 (i.e., the proportion of patches that pass all golden tests) metric. Following prior work [Zhang et al., 2024a], we report three additional reference metrics: (1) *Oracle*, representing the best-case performance where the correct patch is always selected if present among the N candidates; (2) *Adversary*, representing the worst-case performance where an issue is only considered solved if all N candidates are correct; and (3) *Average*, representing the expected performance of randomly selecting one patch from the N candidates.

First, we observe that as the ensemble size increases, the performance of *Average* remains relatively stable, while the *Oracle* value consistently improves. Notably, when the ensemble size reaches 10, *Oracle* achieves an average improvement of 20.80% over *Average* in terms of Pass@1. This observation strongly supports previous findings that although the overall resolution rate remains stable across repeated runs, the set of issues successfully resolved varies across executions [Zhang et al., 2024a, Li et al., 2025, Chen and Flaherty, 2025]. This finding reflects the high diversity in the candidate patches generated by LLMs. *It motivates the significant potential of ensemble reasoning techniques as a promising direction for enhancing LLM-based software issue resolution.*

Moreover, we observe that as the ensemble size increases, the *Adversary* value consistently declines. Specifically, when the ensemble size reaches 10, *Adversary* reduces by 27.01% on average compared to the *Average* in terms of Pass@1. This indicates that although the theoretical upper bound of
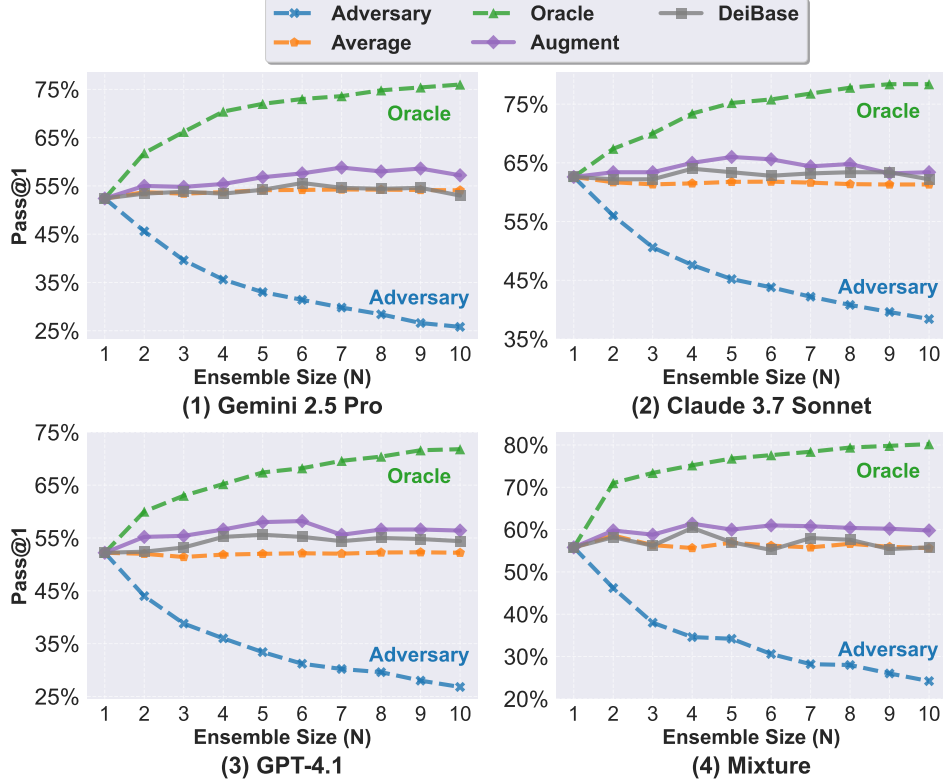
Figure 1: Influence of ensemble size on the effectiveness of existing ensemble reasoning techniques in terms of Pass@1 ($\uparrow$)

ensemble performance (i.e., *Oracle* = 20.80%) improves as the number of candidate patches increases, the theoretical lower bound (i.e., *Adversary* = 27.01%) deteriorates more significantly. This widening gap underscores the increasing difficulty of ensemble reasoning techniques as ensemble size increases. Indeed, as shown in Figure 1, the performance of Augment and DeiBase initially improves with larger ensemble sizes but subsequently degrades. *It motivates the necessity of our patch pruning component, which eliminates redundant and faulty patches, thereby reducing the ensemble space while preserving promising candidates.*

Additionally, on average, both Augment and DeiBase outperform the random selection strategy (i.e., *Average*) by 3.17% and 1.06% in terms of Pass@1, respectively. These results also confirm the potential of ensemble reasoning techniques to enhance the LLM performance in software issue resolution. However, in comparison, *Oracle* significantly achieves the average improvement of 29.49% and 36.29% over Augment and DeiBase, respectively. This substantial performance gap underscores the limitations of existing ensemble reasoning methods in accurately selecting the correct patch and highlights significant potential for further improvement. *It motivates the necessity of designing a more effective patch selection component to further improve the LLM performance in software issue resolution.*

3 Approach

3.1 Overview

Problem setup: Given a software codebase C and a natural-language-described software issue I (which may involve a bug fix or feature implementation), the goal of an ensemble reasoning approach is to automatically generate a diverse set of candidate patches $P = \{p_1, p_2, ..., p_N\}$ and then select a consensus patch p_i ($1 \leq i \leq N$) to resolve I. Applying p_i to the original codebase C yields a
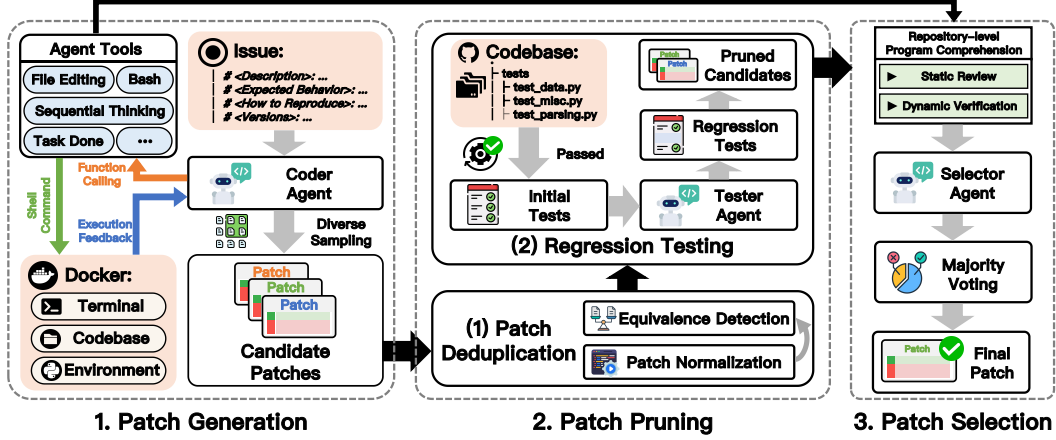
Figure 2: Overview of Trae Agent

modified codebase C'. The selected patch p_i is considered correct if C' passes all golden tests T and satisfies all requirements specified in I.

In this paper, we propose a novel ensemble reasoning approach, `Trae Agent`, designed to enhance the performance of LLMs in software issue resolution. To the best of our knowledge, `Trae Agent` is the first agent-based ensemble reasoning approach for repository-level issue resolution. Due to its generalizable and modular design, `Trae Agent` potentially provides a solid foundation for advancing ensemble reasoning in other complex software engineering tasks. Figure 2 illustrates the overall architecture of `Trae Agent`, which comprises three main components:

- **Patch generation component**(Section 3.2) employs a novel coder agent to generate diverse candidate patches in parallel to improve the ensemble diversity;
- **Patch pruning component** (Section 3.3) performs a hierarchical patch pruning that combines patch deduplication and regression testing strategies to eliminate redundant or faulty patches, thereby reducing the ensemble space while preserving promising candidates;
- **Patch selection component** (Section 3.4) employs a selector agent that simulates a real-world program comprehension process [Eisenbarth et al., 2001, Di Lucca and Di Penta, 2005], constructing repository-level understanding by combining static review with dynamic verification, and ultimately selecting the correct patch through a majority voting strategy.

In the following, we provide a detailed description of each component in `Trae Agent`.

3.2 Patch Generation

Generate candidate patch. As illustrated in Figure 2, this component utilizes a coder agent to generate diverse candidate patches in parallel for each given software issue. Specifically, we implement an effective LLM-based coder agent tailored to address complex GitHub issues. We provide with `Trae Agent` the rich tool ecosystem:

- **File Editing Tool** allows the agent to view files and directories, and make changes (create and edit) to existing files.
- **Bash Tool** provides a persistent command execution interface for the agent to interact with the system, capture output and errors.
- **Sequential Thinking Tool** provides structured problem-solving and analysis capabilities by breaking down complicated problems, iteratively thinking with revision, and hypothesis generation and verification.
- **Task Done Tool** issues a signal of task completion and provides final results and summaries.

As shown in Figure 3, the tools module provides a comprehensive set of capabilities that the agent can use to interact with the environment. Specifically, the coder agent interacts with these tools through
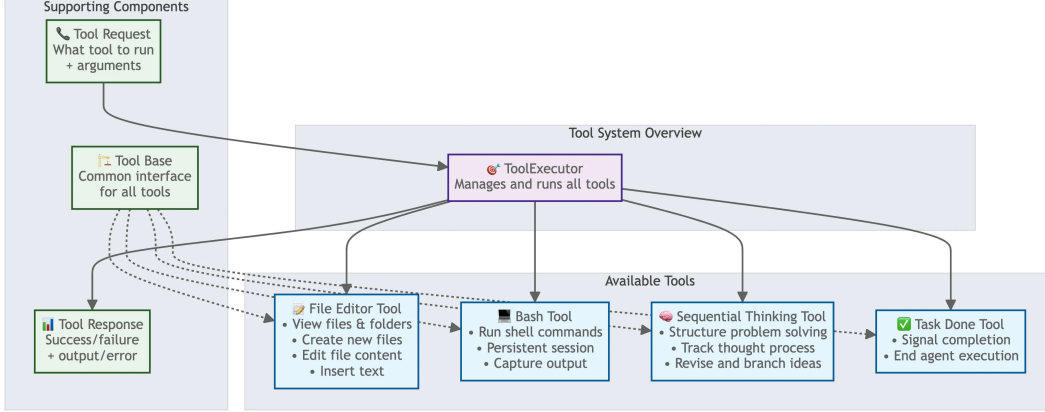
Figure 3: Overview of Tool System

function calling [Abdelaziz et al., 2024, Liu et al., 2025], which are automatically translated into concrete shell commands and executed within a tailored Docker-based environment. The execution results are returned in JSON format, allowing the agent to analyze the execution feedback and refine its understanding. It is worth noting that these tools can be further refined or extended to enhance the capabilities of `Trae Agent` even more effectively. As shown in Figure 4, this coder agent follows a standardized multi-step workflow: (1) analyzing the issue description to understand the target problem; (2) exploring the codebase to locate files relevant to the issue; (3) reproducing the bug to validate its manifestation; (4) diagnosing the root cause through code inspection; (5) generating a code patch to resolve the identified bug; (6) rerunning reproduction tests to verify patch correctness; and (7) summarizing the working process, emulating a realistic commit message. Subsequently, `Trae Agent` leverages the coder agent to generate the candidate patch. Further implementation details are available at our project homepage [Trae-Agent, 2025].

Improve ensemble diversity. To further enhance the diversity of generated candidate patches, `Trae Agent` employs a high-temperature sampling strategy [Zhu et al., 2024, Renze, 2024] across multiple independent runs of the coder agent, thereby further improving the ensemble diversity. As shown in Figure 5, the LLM Client system provides a unified interface for multiple AI providers, works with OpenAI, Anthropic, and Azure and is easy to extend to new providers. In our experiments, the coder agent is instantiated with three state-of-the-art LLMs (Gemini 2.5 Pro, Claude 3.7 Sonnet, and GPT-4.1). Furthermore, we introduce a Mixture setting, in which the three LLMs are used in a round-robin fashion to generate patches, further increasing the diversity of candidates. The generation process terminates once the number of generated candidate patches reaches a pre-defined ensemble size N. This ensemble size is a critical hyper-parameter in all ensemble reasoning techniques, and its impact on performance is discussed in Section 5.2. Overall, this patch generation component enables `Trae Agent` to produce diverse candidate patches, whose union consistently outperforms any individual generation.

Moreover, as shown in Figure 6, to ensure complete observability, the trajectory recording system captures detailed information for each action, including LLM interactions, agent steps, metadata, and error tracking. In addition, we provide a rich terminal output with real-time updates. By default, a lakeview mode is enabled, which asynchronously leverages another LLM to generate concise summarization of agent steps, enhancing the overall user experience.

3.3 Patch Pruning

After generating candidate patches, `Trae Agent` performs a hierarchical patch pruning to reduce the ensemble space, thereby enhancing the effectiveness of the subsequent patch selection. As discussed in Section 2, increasing the ensemble size raises the theoretical upper bound of performance but also lowers the lower bound more sharply, making accurate patch selection more challenging. Existing prompting-based ensemble methods struggle with the optimal solution search in large ensemble spaces. Particularly, our experiments show that, on average, 40% of the candidate patches generated by LLMs are either redundant or incorrect. To mitigate the difficulty of optimal solution

<Codebase Path>
<Github Issue Description>
<Agent Tools>
You are an expert AI software engineering agent. Your primary goal is to resolve a given GitHub issue by navigating the provided codebase, identifying the root cause of the bug, implementing a robust fix, and ensuring your changes are safe and well-tested. Follow these steps methodically:

1. Understand the Problem:
  (1) Begin by carefully reading the user's problem description to fully grasp the issue.
  (2) Identify the core components and expected behavior.

2. Explore and Locate:
  (1) Use the available tools to explore the codebase.
  (2) Locate the most relevant files (source code, tests, examples) related to the bug report.

3. Reproduce the Bug (Crucial Step):
  (1) Before making any changes, you **must** create a script or a test case that reliably reproduces the bug. This will be your baseline for verification.
  (2) Analyze the output of your reproduction script to confirm your understanding of the bug's manifestation.

4. Debug and Diagnose:
  (1) Inspect the relevant code sections you identified.
  (2) If necessary, create debugging scripts with print statements or use other methods to trace the execution flow and pinpoint the exact root cause of the bug.

5. Develop and Implement a Fix:
  (1) Once you have identified the root cause, develop a precise and targeted code modification to fix it.
  (2) Use the provided file editing tools to apply your patch. Aim for minimal, clean changes.

6. Verify and Test Rigorously:
  (1) Verify the Fix: Run your initial reproduction script to confirm that the bug is resolved.
  (2) Prevent Regressions: Execute the existing test suite for the modified files and related components to ensure your fix has not introduced any new bugs.
  (3) Write New Tests: Create new, specific test cases (e.g., using `pytest`) that cover the original bug scenario. This is essential to prevent the bug from recurring in the future. Add these tests to the codebase.
  (4) Consider Edge Cases: Think about and test potential edge cases related to your changes.

7. Summarize Your Work:
  (1) Conclude your trajectory with a clear and concise summary. Explain the nature of the bug, the logic of your fix, and the steps you took to verify its correctness and safety.

Guiding Principle:
Act like a senior software engineer. Prioritize correctness, safety, and high-quality, test-driven development. If you are sure the issue has been solved, you should call the `task_done` to finish the task.

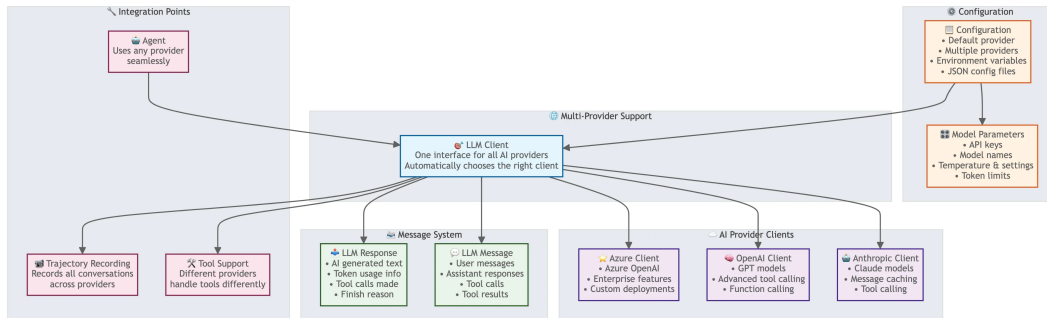Figure 4: Prompt template for the coder agent

Figure 5: Overview of LLM Client System

search, we design a hierarchical pruning method that combines patch deduplication and regression testing strategies to eliminate redundant or faulty patches, thereby reducing the ensemble space while preserving promising candidates. Our experimental results (Section 5) confirm that both patch deduplication and regression testing strategies significantly enhance the performance of `Trae Agent`, and benefit other ensemble reasoning techniques as well. Notably, we first introduce a novel patch pruning strategy for ensemble reasoning, opening a promising direction for future research. The
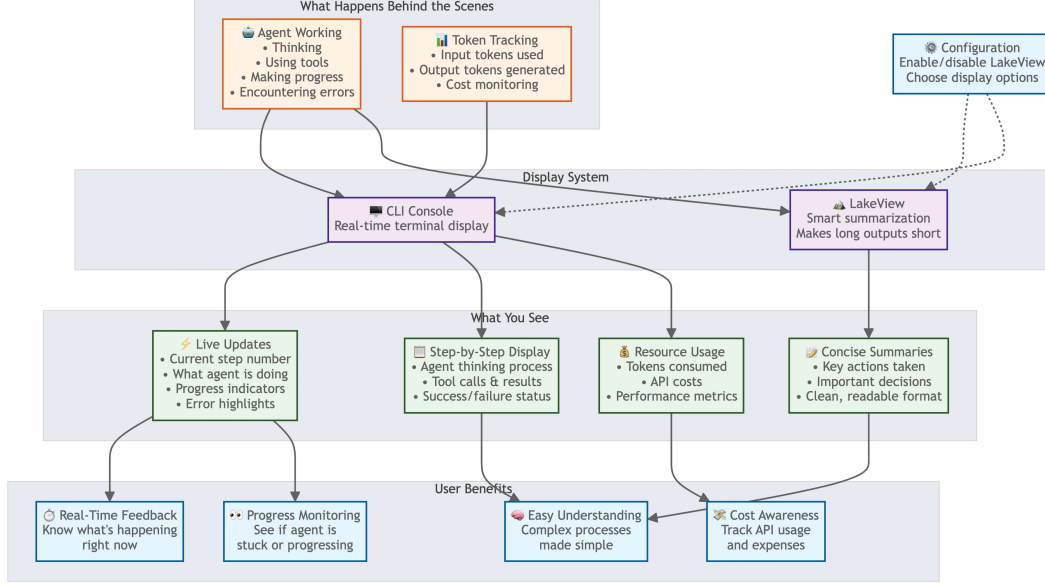
Figure 6: Overview of Trajectory Recording System

following sections present the design and implementation of patch deduplication and regression testing strategies.

Perform patch deduplication. The primary objective of the patch deduplication strategy is to eliminate redundant candidate patches, thereby reducing the ensemble space. Inspired by prior work on equivalent mutant detection [Papadakis et al., 2015, Kintis et al., 2017], we eliminate redundant candidates based on patch normalization and equivalence detection. Specifically, we first implement a patch parser leveraging the Python `unidiff` [Bordese, 2025] package to convert raw patches into structured representations, facilitating precise and reliable patch normalization. We then perform patch normalization to remove semantically irrelevant elements (such as extra spaces, line breaks, and comments) without altering program behavior, yielding a normalized form for each patch. In particular, patches that fail to parse due to syntax errors are deemed invalid and discarded. For subsequent equivalence detection, candidate patches that produce identical normalized representations are considered semantically equivalent, and redundant patches are eliminated accordingly. We acknowledge that, due to the undecidability of program equivalence [Titcheu Chekam et al., 2020, Kim et al., 2022], not all redundant patches can be reliably identified and removed. However, our comprehensive empirical evaluation (Sections 5.3 and 5.4) demonstrates that our proposed strategy effectively reduces redundancy by an average of 28.90% and significantly enhances the practical performance of ensemble reasoning techniques. Furthermore, we will discuss potential enhancements to the patch deduplication process in Section 6.3 to further enhance overall effectiveness.

Perform regression testing. The primary objective of the regression testing strategy is to eliminate faulty candidate patches, thereby further reducing the ensemble space. Specifically, `Trae Agent` employs a tester agent that automatically retrieves and executes regression tests from the original codebase to eliminate candidate patches that fail to preserve existing functionality. Following prior work [Xia et al., 2025], the tester agent first executes all tests in the original codebase and retains those that pass as the initial regression tests. However, not all passing tests are classified as regression tests, as issue resolution may justifiably alter certain existing functionalities, potentially causing some tests to fail. To address this, the tester agent further refines the initial tests by prompting the LLM to identify a subset most likely to represent true regression tests, which are then designated as the final regression tests. Each candidate patch is subsequently applied to the original codebase, and the tester agent executes the final regression tests individually on each patch version. Patches that fail any regression test are discarded, and only those that pass all tests proceed to the patch selection stage. If all candidate patches for an issue fail the regression tests, we conservatively retain the entire set to avoid prematurely discarding potentially valid patches. We acknowledge that the regression tests extracted from the original codebase may contain inaccuracies, potentially introducing false positives

8

Figure 7: Prompt template for the selector agent

during this pruning process. Nonetheless, our empirical study demonstrates that the regression testing strategy effectively eliminates faulty patches, achieving a low error rate of only 3.69%, and substantially enhances the practical performance of ensemble reasoning techniques. Furthermore, the impact of inaccurate regression tests is analyzed in detail in Section 6.2.

3.4   Patch Selection

The primary goal of the patch selection component is to accurately identify the correct patch from the pruned candidates. Real-world issues and fixes frequently span multiple files and modules, necessitating cross-file reasoning, contextual awareness, and verification of patch correctness within the broader codebase. Therefore, effective patch selection requires an accurate repository-level understanding. However, existing prompting-based ensemble reasoning approaches are typically stateless and single-turn, lacking persistent memory and tool integration. These limitations hinder their ability to track dependencies, perform validation steps, and build a coherent, global understanding of the software repository.

Obtain repository-level understanding. To address the challenge of repository-level understanding, `Trae Agent` introduces a selector agent that simulates a real-world program comprehension process [Eisenbarth et al., 2001, Di Lucca and Di Penta, 2005]. Figure 7 shows the detailed prompt template for the selector agent. It iteratively gathers and analyzes relevant code snippets (such as those referenced in the issue description, modified by patches, or related through dependencies) to build the static understanding, and collects execution traces from automatically generated tests to build the dynamic understanding. To achieve this, we also equip the selector agent with tools (introduced in Section 3.2) that provide access to the complete codebase and execution environment for each issue. In addition to gathering and analyzing relevant code snippets for static understanding, the selector agent leverages these tools to generate and execute unit tests, thereby facilitating dynamic understanding, as illustrated in Figure 8. The selector agent iteratively leverages these tools until it obtains sufficient repository-level understanding to make a final patch selection. Additionally, to control computational overhead and prevent excessive token usage, we impose an upper bound of 30 interaction rounds between the selector agent and the agent tools.

Implement majority voting. To further mitigate the risk of LLM hallucinations and enhance solution consistency, we implement a majority voting strategy into the selector agent. Specifically, given N candidate patches, the selector agent is executed in parallel for N iterations, and the patch receiving
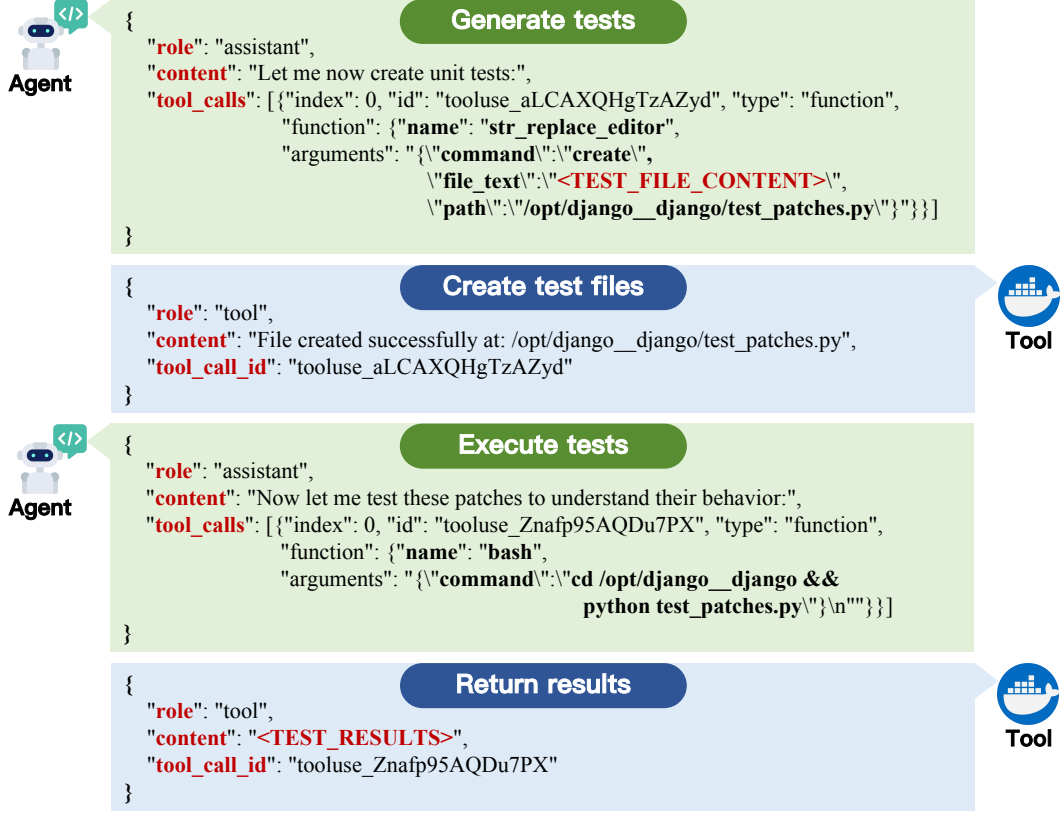
9

Figure 8: An example of Trae Agent to execute generated tests

the highest number of votes is selected as the final output. To improve efficiency, if the first $\lceil \frac{N}{2} \rceil$ votes are unanimous, this consensus patch is immediately returned as the final result, and the remaining $(N - \lceil \frac{N}{2} \rceil)$ executions are skipped. In cases where the N votes are evenly distributed across multiple candidate patches, potentially indicating the presence of several correct solutions, `Trae Agent` randomly selects one from among the top-voted candidates. The impact of this majority voting strategy is further analyzed in Section 5.3. Through this patch selection component, `Trae Agent` can effectively select the most plausible patch to improve LLM-based issue resolution.

4   Evaluation Design

Our study aims to address the following research questions (RQs):

- **RQ1**: How does `Trae Agent` perform in improving the issue resolution performance of LLMs compared to the state-of-the-art techniques?

- **RQ2**: How do hyper-parameters affect `Trae Agent`'s effectiveness?

- **RQ3**: How does each main component in `Trae Agent` contribute to the overall effectiveness?

- **RQ4**: How does the ensemble space affect the effectiveness of patch selection?

4.1   Benchmarks

We evaluate the effectiveness of `Trae Agent` on SWE-bench [Jimenez et al., 2023], a widely-used benchmark for automated software issue resolution that comprises 2,294 real-world GitHub issues. However, prior studies [OpenAI, 2024, Wang et al., 2025c] have shown that the full SWE-bench dataset contains a significant number of low-quality and noisy instances. To address this limitation, OpenAI [OpenAI, 2025b] curated a high-quality subset called SWE-bench Verified [OpenAI, 2024],

which includes 500 GitHub issues that have been manually verified by professional software developers to be reliable and non-problematic. Following existing work [Xia et al., 2025, Wang et al., 2025c, Zhang et al., 2025], our evaluation focuses on this verified version. For each issue, all evaluated techniques receive only the issue description and the original codebase as input. SWE-bench also provides some golden tests for each GitHub issue to evaluate whether the generated patch solves the corresponding issue. Note that all regression tests used by `Trae Agent` are part of the original codebase, ensuring a fair and realistic evaluation setting.

4.2 Metrics

We use **Pass@1** to assess the effectiveness of the studied ensemble reasoning techniques. It measures the functionality correctness of a generated patch by determining whether it successfully resolves the given software issue. For each issue, SWE-bench provides golden tests to verify patch correctness. In our evaluation, the ensemble reasoning technique selects a single patch per issue; if this patch passes all associated tests in the SWE-bench, the issue is deemed successfully resolved. It is worth noting that Pass@1 is a stringent evaluation criterion, making improvements to this metric both technically challenging and practically meaningful [Xia et al., 2025, Ruan et al., 2025]. A higher Pass@1 score indicates better effectiveness.

4.3 Compared Techniques

To comprehensively evaluate the performance of `Trae Agent`, we compare it with state-of-the-art ensemble reasoning techniques:

- **Augment** [Code, 2025]: implements a prompting-based ensemble technique inspired by the LLM-as-a-judge paradigm. It prompts the LLM to compare each candidate against the issue description and choose the best match.
- **DeiBase** [Zhang et al., 2024a]: prompts the LLM to generate the detailed justifications and assign the confidence scores to each candidate patch, thereby selecting the one with the highest score as the final patch.

To further demonstrate the effectiveness and generality of the patch pruning component in our `Trae Agent`, we integrate this component into both ensemble reasoning techniques, resulting in two enhanced baselines: **Augment w/ Pruning** and **DeiBase w/ Pruning**. To ensure a fair comparison, all ensemble reasoning techniques are evaluated using the same set of candidate patches generated by our patch generation component.

4.4 Implementation Details

To evaluate the performance of `Trae Agent`, we employ three state-of-the-art LLMs: Gemini 2.5 Pro [DeepMind, 2025] (version `gemini-2.5-pro-preview-06-05`), Claude 3.7 Sonnet [Anthropic, 2025b] (version `claude-3-7-sonnet-20250219`), and GPT-4.1 [OpenAI, 2025a] (version `gpt-4.1-2025-04-14`). These models are used for patch generation. Among them, Claude 3.7 Sonnet demonstrated the best performance (shown in Figure 1); therefore, for consistency and fair comparison, we adopt Claude 3.7 Sonnet as the base model for all ensemble reasoning techniques. Specifically, we access Gemini 2.5 Pro, Claude 3.7 Sonnet, and GPT-4.1 through the official APIs provided by Google AI [AI, 2025], Anthropic [Anthropic, 2025a], and OpenAI [OpenAI, 2025b], respectively. All other parameters are kept at their default values. To evaluate the functionality correctness of the selected patches, we use the official SWE-bench Docker environment to ensure consistency and reliability in testing [Jimenez et al., 2023].

5 Results and Analysis

5.1 RQ1: Overall Effectiveness of Trae Agent

5.1.1 Process:

To answer RQ1, we evaluate the performance of `Trae Agent` alongside four state-of-the-art ensemble reasoning baselines (Augment, Augment w/ Pruning, DeiBase, and DeiBase w/ Pruning) across

Table 1: Effectiveness comparison in terms of Pass@1 (↑).

Technique	Gemini 2.5 Pro	Claude 3.7 Sonnet	GPT-4.1	Mixture
Adversary	39.60%	50.60%	38.80%	38.00%
Average	53.40%	61.33%	51.40%	56.33%
Oracle	66.20%	70.00%	63.00%	73.40%
Augment	55.40%±0.60%	63.13%±0.31%	54.87%±0.92%	58.93%±0.23%
Augment w/ Pruning	59.27%±0.31%	64.33%±0.42%	56.60%±0.53%	61.07%±0.31%
DeiBase	53.53%±0.23%	62.33%±0.42%	53.13%±**0.12%**	55.87%±0.76%
DeiBase w/ Pruning	57.33%±0.12%	63.87%±0.31%	56.00%±0.35%	58.07%±1.01%
Trae Agent	**62.27%**±**0.12%**	**66.40%**±**0.20%**	**59.00%**±0.20%	**65.67%**±**0.23%**
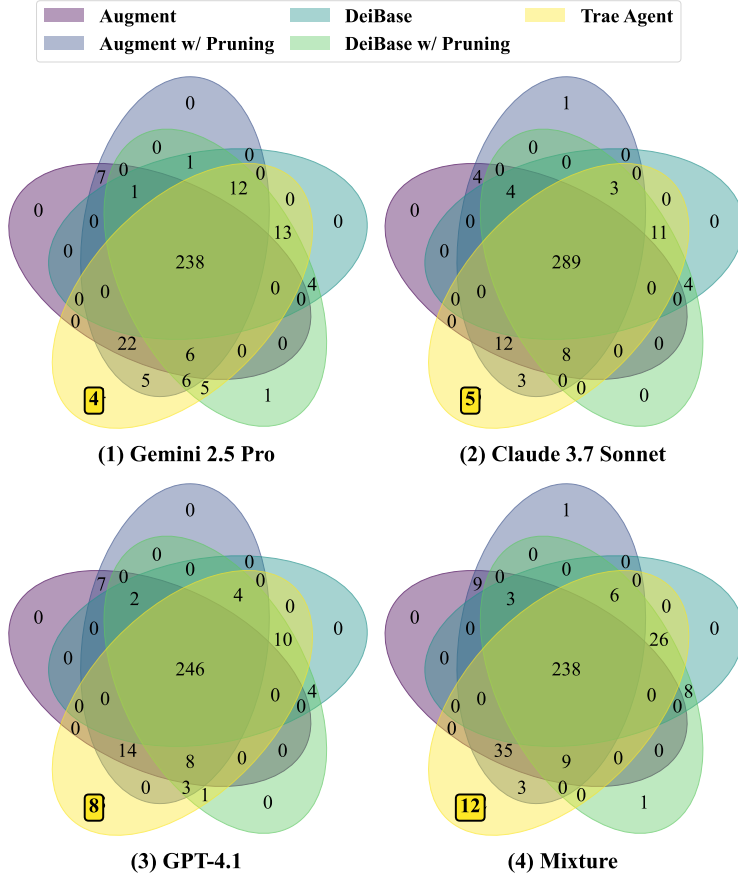
Figure 9: Number of uniquely resolved issues across five studied ensemble reasoning techniques on SWE-bench Verified

three leading LLMs (Gemini 2.5 Pro, Claude 3.7 Sonnet, and GPT-4.1). Additionally, we introduce a Mixture setting, where the three LLMs generate patches in a round-robin manner to further enhance the diversity of the candidate patches. The effectiveness of each technique is assessed on the widely-used SWE-bench Verified benchmark using two metrics (Pass@1 and the number of uniquely resolved issues). Moreover, we include three reference baselines (introduced in Section 2): *Oracle* (representing the best-case performance of an ensemble technique), *Adversary* (representing the worst-case performance of an ensemble technique), and *Average* (representing the expected performance of a random baseline). The ensemble size N is set to 3. To further mitigate the impact of randomness, each experiment is repeated three times.

5.1.2 Results:

Table 1 presents a comprehensive comparison of all studied ensemble reasoning techniques in terms of effectiveness. First, we observe that all five ensemble reasoning techniques consistently outperform the *Average* baseline in terms of Pass@1, validating the effectiveness of ensemble reasoning techniques and reinforcing the motivation for our design of `Trae Agent`. Notably, `Trae Agent` achieves the best performance across all ensemble reasoning techniques. Specifically, it demonstrates an improvement of 5.01%~12.86% compared to the baselines in terms of Pass@1. Furthermore, the *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) confirms that all p-values are lower than 8.00×10^{-6}, demonstrating the statistically significant superiority of `Trae Agent` over all baselines in terms of Pass@1. In addition, `Trae Agent` exhibits strong stability, with an average standard deviation of only 0.19%, lower than those of other baselines (0.38%~0.52%), thereby further reducing the impact of randomness. Moreover, as shown in Figure 9, `Trae Agent` also achieves the highest number of uniquely resolved issues, further demonstrating its superior effectiveness compared to ensemble reasoning baselines.

Secondly, Table 1 also shows that both Augment w/ Pruning and DeiBase w/ Pruning consistently outperform their original counterparts (Augment and DeiBase) in terms of Pass@1. On average, the patch pruning component improves the performance of Augment and DeiBase by 3.91% and 4.72% in terms of Pass@1, respectively. This result empirically validates the effectiveness of our patch pruning component. Furthermore, the *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) confirms that all p-values are lower than 1.30×10^{-5}, which statistically demonstrates that the patch pruning component provides a substantial enhancement to the ensemble reasoning baselines in terms of Pass@1.

5.2 RQ2: Influence of Hyper-parameter

5.2.1 Setup:

The ensemble size (N) is a critical hyper-parameter in ensemble reasoning techniques. In this research question, we investigate the impact of the ensemble size on the performance of `Trae Agent` and other compared ensemble reasoning techniques. Specifically, for $1 \leq N \leq 10$, we evaluate and compare the effectiveness of `Trae Agent` alongside four baselines in terms of Pass@1.

5.2.2 Results:

Figure 10 illustrates the performance trends of different ensemble reasoning techniques as the ensemble size increases, evaluated using the Pass@1 metric. Across all ensemble size settings, `Trae Agent` consistently outperforms all four baselines, achieving an average improvement of 5.83%~14.60% in terms of Pass@1. These results demonstrate the stable effectiveness of `Trae Agent` under varying ensemble size settings. Furthermore, the *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) confirms that all p-values are smaller than 3.74×10^{-12}, validating the statistically significant superiority of `Trae Agent` across different ensemble sizes.

In addition, as the ensemble size increases, the performance of the four baselines generally improves initially but subsequently declines. This trend can be attributed to two key factors: (1) while the theoretical upper bound increases with larger ensemble sizes, the lower bound degrades more rapidly (as discussed in Section 2); and (2) the increasing context length required to process more candidate patches leads to dilution of relevant information, making accurate patch selection more difficult for LLMs. In contrast, `Trae Agent` consistently demonstrates an upward trajectory in terms of Pass@1 with increasing ensemble size, demonstrating superior effectiveness and scalability. We hypothesize that further scaling may continue to enhance `Trae Agent`'s effectiveness (see Section 6.3); however, this entails a trade-off between effectiveness and computational cost.

Additionally, we observe that Augment w/ Pruning and DeiBase w/ Pruning consistently outperform their original counterparts (Augment and DeiBase) across all ensemble size settings in terms of Pass@1. This result further substantiates the effectiveness of the patch pruning component. On average, the patch pruning component yields improvements of 3.91% and 3.74% for Augment and DeiBase, respectively. Furthermore, the *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) confirms that all p-values are smaller than 1.95×10^{-11}, indicating the patch pruning component can significantly enhance existing ensemble reasoning techniques.
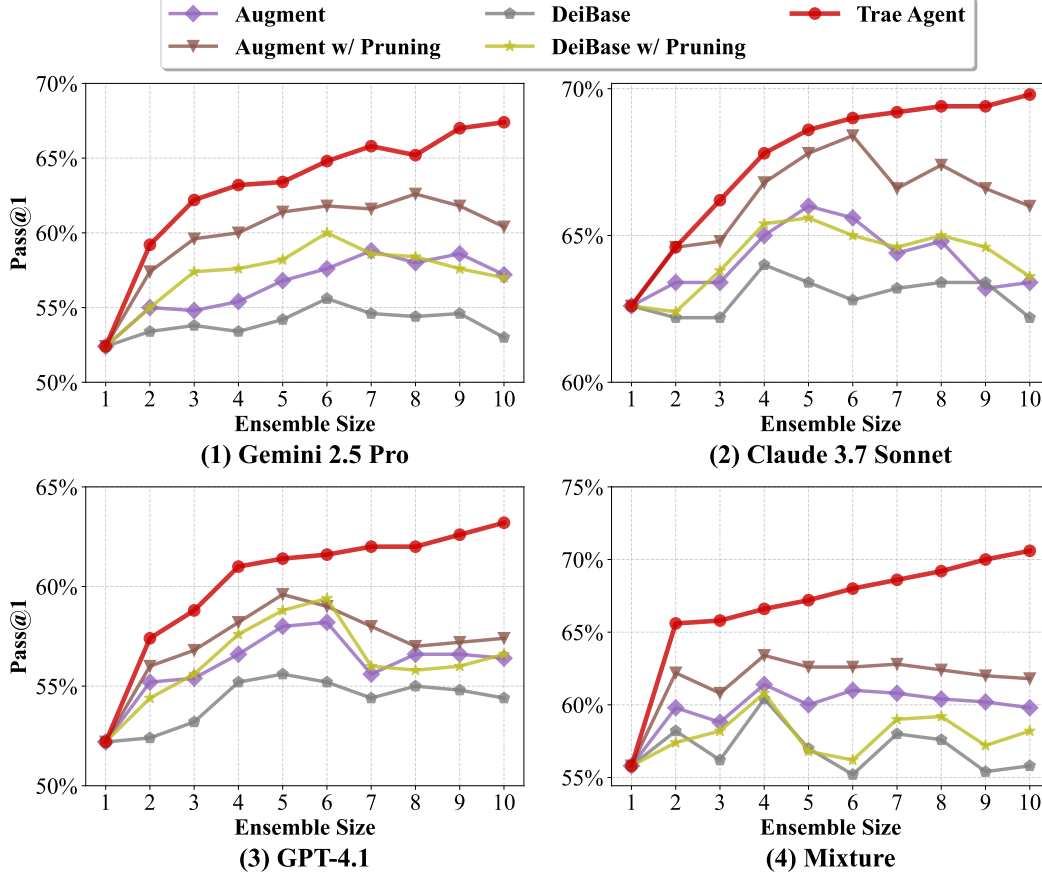
Figure 10: Influence of the ensemble size in terms of Pass@1 (↑)

Table 2: Comparison between `Trae Agent` and its five ablation variants in terms of Pass@1 (↑).

Technique	Gemini 2.5 Pro	Claude 3.7 Sonnet	GPT-4.1	Mixture
Trae Agent$_{woD}$	60.00%	64.00%	57.00%	63.20%
Trae Agent$_{woR}$	60.20%	64.60%	56.80%	63.40%
Trae Agent$_{woP}$	58.40%	63.60%	56.20%	61.80%
Trae Agent$_A$	59.80%	64.60%	57.20%	61.80%
Trae Agent$_{woM}$	59.60%	63.60%	57.40%	62.60%
Trae Agent	62.27%	66.40%	59.00%	65.67%

5.3 RQ3: Contribution of Main Components

5.3.1 Variants:

To assess the contributions of the main components in `Trae Agent`, we construct and evaluate five ablation variants. For the patch pruning component, which comprises both patch deduplication and regression testing strategies, we design three variants: (1) `Trae Agent`$_{woD}$, which removes the patch deduplication strategy; (2) `Trae Agent`$_{woR}$, which removes the regression testing strategy; and (3) `Trae Agent`$_{woP}$, which removes the entire patch pruning component. For the patch selection component, we construct two variants: (4) `Trae Agent`$_A$, which replaces the selector agent in `Trae Agent` with the advanced prompting-based Augment, allowing us to evaluate the effectiveness of the selector agent; (5) `Trae Agent`$_{woM}$, which removes the majority voting strategy to assess its contribution within the patch selection component.

Table 3: Correlation coefficients between ensemble space and selection effectiveness. A correlation coefficient with an absolute value closer to 1.0 indicates a stronger correlation.

Correlation	Gemini 2.5 Pro	Claude 3.7 Sonnet	GPT-4.1	Mixture
Pearson's r	0.91 (*very strong*)	0.73 (*strong*)	0.78 (*strong*)	0.89 (*very strong*)
Spearman's ρ	1.00 (*perfect*)	0.80 (*very strong*)	0.80 (*very strong*)	1.00 (*perfect*)
Kendall's τ	1.00 (*perfect*)	0.67 (*strong*)	0.67 (*strong*)	1.00 (*perfect*)

5.3.2 Results:

Table 2 presents the comparison results of `Trae Agent` and its five ablation variants in terms of Pass@1. First, we observe that `Trae Agent` consistently outperforms the three pruning-related variants (`Trae Agent`$_{woP}$, `Trae Agent`$_{woD}$, and `Trae Agent`$_{woR}$). Specifically, `Trae Agent` achieves average improvements of 5.57%, 3.73%, and 3.42% over`Trae Agent`$_{woP}$, `Trae Agent`$_{woD}$, and `Trae Agent`$_{woR}$, respectively, demonstrating the contributions of the patch pruning component and its patch deduplication and regression testing strategies. These results reinforce the critical role of the patch pruning component in enhancing the overall performance of `Trae Agent`. Furthermore, as shown in RQ1 and RQ2, incorporating the patch pruning component into other ensemble reasoning techniques (i.e., Augment w/ Pruning and DeiBase w/ Pruning) also leads to performance improvements, further validating its generalizability and utility.

Second, `Trae Agent` achieves an average improvement of 4.08% over `Trae Agent`$_A$ in terms of Pass@1, validating the effectiveness of the selector agent and further demonstrating its superiority over existing prompting-based ensemble reasoning techniques. In addition, `Trae Agent` outperforms `Trae Agent`$_{woM}$ by an average of 4.14% in terms of Pass@1, highlighting the contribution of the majority voting strategy in enhancing the overall performance. This result confirms the necessity of incorporating voting-based mechanisms to mitigate selection instability and reduce the impact of LLM hallucinations.

Furthermore, the *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) confirms that all p-values are smaller than 7.32×10^{-3}, exhibiting the statistically significant advantage of `Trae Agent` over all variants. Overall, each of the main components contributes substantially to the overall effectiveness of `Trae Agent`.

5.4 RQ4: Influence of Ensemble Space

5.4.1 Process

To address RQ4, we first examine the impact of `Trae Agent` and its three pruning-related variants (i.e., `Trae Agent`$_{woP}$, `Trae Agent`$_{woD}$, and `Trae Agent`$_{woR}$ introduced in RQ3) on the ensemble space (i.e., the number of remaining candidate patches after patch pruning) and then conduct a in-depth correlation analysis between ensemble space and selection effectiveness (measured by Pass@1). Specifically, we fix the ensemble size to 10 and evaluate four techniques: `Trae Agent`$_{woP}$ (which removes the entire patch pruning component including both patch deduplication and regression testing strategies), `Trae Agent`$_{woR}$ (which removes the regression testing strategy while retaining the patch deduplication strategy), `Trae Agent`$_{woD}$ (which removes the patch deduplication strategy while retaining the regression testing strategy), and `Trae Agent` (which includes the complete patch pruning component).

To further investigate the correlation between the ensemble space and effectiveness, we perform a statistical correlation analysis [Zhang and Mesbah, 2015, Li and Zhang, 2017, Shin et al., 2019] using three widely-adopted correlation coefficients: *Pearson's* r coefficient [Cohen et al., 2009] (that measures linear correlation), *Spearman's* ρ coefficient [Spearman, 1961] (that captures monotonic relationships and is closely related to Pearson's r coefficient), and *Kendall's* τ coefficient [Kendall, 1938] (that measures rank correlation). According to established thresholds in prior work [Buse and Weimer, 2009, Sebastian et al., 2017], a coefficient with an absolute value of 1.0 denotes *perfect* correlation; values in the range [0.8,1.0) indicate *very strong* correlation, [0.6,0.8) *strong* correlation, [0.4,0.6) *moderate* correlation, [0.2,0.4) *weak* correlation, and (0,0.2) *very weak* correlation. A value of 0 indicates *no correlation*.
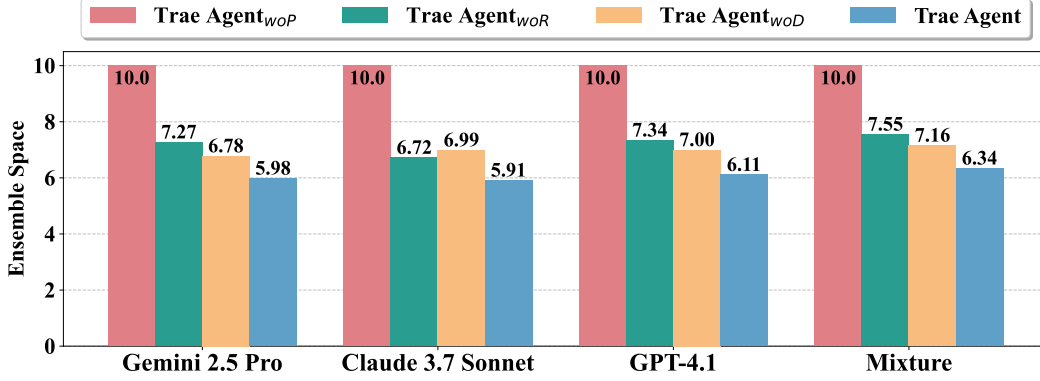
Figure 11: Influence of the patch pruning component and its patch deduplication and regression testing strategies in terms of ensemble space ($\downarrow$)

Model	% Resolved	Org	Date	Logs	Trajs	Site
🆕 TRAE	75.20	▬	2025-06-12	✓	✓	🔗
🆕 Refact.ai Agent	74.40	⟦	2025-06-03	✓	✓	🔗
🆕✅ Moatless Tools + Claude 4 Sonnet	70.80	-	2025-06-11	✓	✓	🔗
🆕 Refact.ai Agent	70.40	⟦	2025-05-15	✓	✓	🔗
🆕✅ OpenHands + Claude 4 Sonnet	70.40	👓	2025-05-24	✓	✓	🔗
🆕 Augment Agent v1	70.40	🔧	2025-06-10	✓	✓	🔗
🆕✅ SWE-agent + Claude 4 Sonnet	66.60	🗃	2025-05-22	✓	✓	🔗
✅ OpenHands	65.80	👓	2025-04-15	✓	✓	🔗
Augment Agent v0	65.40	🔧	2025-03-16	✓	✓	🔗
🆕 PatchPilot-v1.1	64.60	⚪	2025-05-03	✓	✓	-
✅ SWE-agent + Claude 3.7 Sonnet w/ Review Heavy	62.40	🗃	2025-02-25	✓	✓	🔗

Figure 12: Snapshot of the official SWE-bench Verified leaderboard as of July 2025.

5.4.2 Results

As shown in Figure 11, `Trae Agent`, `Trae Agent`$_{woR}$, and `Trae Agent`$_{woD}$ all reduce the ensemble space compared to `Trae Agent`$_{woP}$. Specifically, they achieve average reductions of 27.80%, 30.22%, and 39.15%, respectively. These results demonstrate that the patch pruning component, along with its patch deduplication and regression testing strategies, effectively reduces the ensemble space. Combined with the findings in Table 2 (as discussed in RQ3), which show that these components contribute to improved selection effectiveness, we can qualitatively conclude that reducing the ensemble space is beneficial to enhance patch selection performance.

To further quantitatively assess the correlation between the ensemble space and selection effectiveness, we compute the *Pearson's r*, *Spearman's ρ*, and *Kendall's τ* between the two metrics (ensemble space and Pass@1). As shown in Table 3, the absolute values of *Pearson's r* range from 0.73 to 0.91, *Spearman's ρ* range from 0.80 to 1.00, and *Kendall's τ* range from 0.67 to 1.00, respectively. These results consistently indicate a strong correlation across all three correlation metrics, suggesting that reducing the ensemble space is closely associated with improved selection effectiveness.

Table 4: Quality of regression tests in terms of Accuracy (↑), Precision (↑), Recall (↑), F1-Score (↑).

Metric	TP	TN	FP	FN	Total
# Instances	10,424	2,231	6,608	737	20,000
Accuracy = (TP+TN)/Total					63.28%
Precision = TP/(TP+FP)					61.20%
Recall = TP/(TP+FN)					93.40%
F1-Score = (2×Precision×Recall)/(Precision+Recall)					73.95%

* TP (true positive): patch passes regression tests and is correct;

* TN (true negative): patch does not pass regression tests and incorrect;

* FP (false positive): patch passes regression tests but is incorrect;

* FN (false negative): patch does not pass regression tests but is correct.

Table 5: Influence of the regression testing strategy in terms of the ratios of all-correct (↑) and all-incorrect (↓) instances.

Metric	Technique	Gemini 2.5 Pro	Claude 3.7 Sonnet	GPT-4.1	Mixture
All-correct	w/o Regression	25.80%	38.40%	26.80%	24.20%
	w/ Regression	40.80%	48.00%	37.00%	37.40%
All-incorrect	w/o Regression	24.00%	21.60%	28.20%	19.80%
	w/ Regression	28.60%	25.80%	32.40%	24.60%

6 Discussion

6.1 Practical Impact

Building autonomous agent-based systems for software engineering tasks has emerged as a leading research direction in recent years. SWE-bench, a widely adopted benchmark, encompasses real-world software engineering tasks such as bug fixing and feature implementation. As of July 2025, SWE-bench has been downloaded over 500,000 times, with more than 90 agent-based techniques submitted to its official leaderboard. The SWE-bench leaderboard [bench Team, 2025] continuously tracks the state-of-the-art progress in agent-based software engineering systems. **As shown in Figure 12, Trae Agent has achieved the first place on the SWE-bench Verified leaderboard, with a notable Pass@1 score of 75.20%.** Furthermore, our open-source GitHub repository has attracted over 8,000 stars (as of July 2025), indicating substantial community interest and adoption. Our findings suggest that integrating ensemble reasoning techniques into existing agent-based systems can significantly improve their ability to resolve complex real-world issues. We believe that this insight could serve as a guiding principle for the design of future agent-based software engineering frameworks.

6.2 Quality of Regression Tests

In the patch pruning component, the regression tests selected by the tester agent play a critical role in determining the overall effectiveness of `Trae Agent`. To evaluate the quality of these regression tests, we assess the consistency between their pruning decisions and the ground-truth results provided by the golden tests in the SWE-bench. Specifically, we examine whether the regression tests accurately determine the correctness of the candidate patches. As shown in Table 4, the regression tests achieve an Accuracy of 63.28%, a Precision of 61.20%, a Recall of 93.40%, and an F1-Score of 73.95%, respectively. We further analyze the two types of regression test mispredictions (i.e., FP and FN), which directly affect the effectiveness of patch pruning. The overall misprediction rate ($\frac{FP+FN}{Total}$) is 36.73%. Among these, FP cases (i.e., incorrect patches mistakenly passing the regression tests) can be viewed as a conservative patch pruning strategy that preserves erroneous patches for subsequent selection. In contrast, FN cases (i.e., correct patches mistakenly discarded) are more detrimental to overall effectiveness, as they reduce the correct candidate patches. In particular, the FP rate is 33.04%, while the FN rate accounts for only 3.69%. This indicates that the risk of discarding correct patches is relatively low.

In addition, we further investigate the impact of regression tests on instance-level correctness. Specifically, each instance is associated with multiple candidate patches. If all candidate patches are correct, the instance is guaranteed to be resolved successfully; conversely, if all candidates are incorrect, the instance cannot be resolved. Therefore, for ensemble reasoning techniques, it is desirable to maximize the number of all-correct instances and minimize the number of all-incorrect instances. To this end, we evaluate the ratios of all-correct and all-incorrect instances before and after applying the regression testing strategy. As shown in Table 5, the regression testing strategy increases the ratio of all-incorrect instances by 4.45% on average, but more notably increases the ratio of all-correct instances by 12.00% on average. These findings further support the effectiveness of the regression testing strategy and provide an in-depth explanation for its positive impact on the overall performance of `Trae Agent`.

6.3   Future Work

`Trae Agent` is the first agent-based ensemble reasoning approach for repository-level issue resolution. While its effectiveness has been demonstrated through comprehensive empirical studies, several aspects of `Trae Agent` can be further enhanced in future work:

(1) **Enhancing ensemble diversity**: Currently, `Trae Agent` relies on a single coder agent with a high-temperature sampling strategy for diverse patch generation. We plan to integrate multiple issue resolution agents (e.g., Agentless [Xia et al., 2025], Moatless [Örwall, 2025], and Open-Hands [Wang et al., 2025b]) to better leverage their complementary strengths. This integration is expected to increase the diversity of candidates, thereby improving the ensemble performance.

(2) **Scaling ensemble size**: Our current experiments (as discussed in RQ2) evaluate ensemble sizes ranging from 1 to 10, with results showing a consistent performance improvement of `Trae Agent` as the ensemble size increases. In future work, we plan to investigate the potential performance gains of larger ensemble sizes, while carefully considering cost-efficiency constraints.

(3) **Enhancing patch deduplication**: Recent studies have explored leveraging LLMs for program equivalence detection, such as identifying equivalent mutants [Ma et al., 2023, Tian et al., 2024]. Building on this, we plan to investigate LLM-based patch deduplication techniques to further enhance the overall performance of `Trae Agent`.

7   Threats and Validity

Construct Validity. This threat mainly arises from the inherent randomness of LLMs. To mitigate this, we conduct a large-scale empirical study and release the replication package for practitioners. Moreover, we ensure consistency by repeating all experiments three times in RQ1. Notably, the standard deviations of Pass@1 for Augment, Augment w/ Pruning, DeiBase, DeiBase w/ Pruning, and `Trae Agent` are only 0.0052, 0.0039, 0.0038, 0.0044, and 0.0019, respectively, indicating high robustness across runs. In addition, a *Wilcoxon Signed-Rank Test* [Wilcoxon et al., 1963] (at a significance level of 0.05) yields p-values exceeding 0.38 for all comparisons, indicating no statistically significant differences across the three repeated experimental results and further strengthening the reliability of our findings.

External Validity. This threat mainly lies in our experimental subjects. To address this, we adopt the widely-used SWE-bench benchmark and multiple evaluation metrics commonly used in the issue resolution domain [Zhang et al., 2024b,a, Ruan et al., 2025, Wang et al., 2025b]. We further evaluate `Trae Agent` against four state-of-the-art ensemble reasoning techniques across three leading LLMs and the widely-used SWE-bench benchmark, ensuring comprehensive comparisons. In future work, we plan to extend our evaluation to additional benchmarks and LLMs to further assess the generalizability and robustness of `Trae Agent` across diverse settings.

8 Related Work

8.1 Automatic Software Issue Resolution

Automatic software issue resolution is a critical task in software engineering and has attracted increasing research interest in recent years. Numerous techniques have been proposed, among which agent-based approaches have gained particular prominence.

OpenDevin [Wang et al., 2025b] (renamed OpenHands) is among the earliest agent-based frameworks, leveraging a planning mechanism based on user requirements and utilizing tools (such as file editors, terminals, and web search engines) to iteratively accomplish complex tasks. SWE-agent [Yang et al., 2024] introduces a custom agent-computer interface, enabling agents to interact with the codebase through operations such as viewing and editing files. Moatless [Örwall, 2025] enhances issue resolution by equipping agents with code search tools and retrieval strategies to identify relevant code locations. AutoCodeRover [Zhang et al., 2024b] further refines the code search capability by representing software projects as abstract syntax trees, allowing agents to effectively retrieve contextual information and locate faults. Building on AutoCodeRover, SpecRover [Ruan et al., 2025] improves the specifications by generating function summaries, and incorporates the generation of reproduction tests to assist in patch generation. Agentless [Xia et al., 2025] adopts a standardized operational pipeline comprising fault localization, patch generation, and patch verification, without requiring agents to dynamically decide on future actions or interact with complex external tools. MarsCode [Liu et al., 2024] Agent combines advanced code analysis techniques (i.e., code knowledge graph) with LLM capabilities to provide a systematic process for fault localization, candidate patch generation, and patch validation.

Unlike existing individual patch generation techniques, `Trae Agent` can be seamlessly integrated with them to build more effective ensemble reasoning systems. Furthermore, due to its generalizable and modular design, `Trae Agent` potentially provides a promising foundation for advancing ensemble reasoning in broader and more complex software engineering tasks.

8.2 Ensemble Techniques for LLM

Ensemble learning [Polikar, 2012, Sagi and Rokach, 2018] has been widely adopted in machine learning and deep learning models to reduce prediction bias and improve generalization by selecting a consensus solution from multiple models, thereby mitigating the limitations inherent to any single model. Traditional ensemble methods (e.g., bagging [Breiman, 1996], boosting [Freund and Schapire, 1997], and stacking [Džeroski and Ženko, 2004]) have achieved notable success across a variety of domains, including image classification [Chen et al., 2019], natural language processing [Zhang and Shafiq, 2024], and anomaly detection [Han et al., 2021].

Recent advances have demonstrated the potential of ensemble techniques to enhance the performance of LLM-based agent systems across a range of tasks. In the mathematical reasoning domain, Snell et al. [2025] conduct a comprehensive study and introduce the Best-of-N approach, which generates multiple solutions in parallel from a base LLM and employs a reward model to select the highest-scoring solution. In the domain of competition-level code generation, Li et al. [2025] propose S*, which leverages both the execution feedback from public and LLM-generated test cases along with a clustering-based selection mechanism to guide the final code selection. In addition, Mahmud et al. [2025] develop EnsLLM, which utilizes CodeBLEU [Ren et al., 2020] for syntactic similarity and CrossHair [CrossHair, 2025] (a property-based testing tool) for behavior similarity. Both similarities are then integrated using a voting mechanism to select the most reliable candidate solution.

In the more complex task of software issue resolution, Chen and Flaherty [2025] propose Augment, an agent that applies the LLM-as-a-judge paradigm to assess the semantic alignment between the given GitHub issue and a set of candidate patches, ultimately selecting the most relevant patch. In addition, Zhang et al. [2024a] introduce DeiBase, which prompts the LLM to generate detailed explanations and confidence scores for each candidate patch, thereby selecting the one with the highest score. In contrast to existing ensemble techniques, we propose `Trae Agent`, the first agent-based ensemble reasoning approach for repository-level software issue resolution. `Trae Agent` is designed to enhance LLM-based issue resolution by formulating it as an optimal solution search problem and enhancing LLM performance in issue resolution through a modular agent-based architecture.

9 Conclusion

In this work, we present `Trae Agent`, the first agent-based ensemble reasoning approach for repository-level issue resolution, which substantially improves the effectiveness of LLMs. `Trae Agent` formulates ensemble reasoning as an optimal solution search problem and tackles two fundamental challenges, i.e., large ensemble spaces and repository-level understanding, through modular agents for generation, pruning, and selection. To comprehensively evaluate its performance, we conduct extensive experiments using three leading LLMs on the widely-used SWE-bench benchmark. The experimental results show that `Trae Agent` consistently outperforms four state-of-the-art ensemble reasoning baselines across all evaluation settings, demonstrating its effectiveness in enhancing LLM-based software issue resolution.

Contributors

Core Contributors
Pengfei Gao[*]
Zhao Tian[*]
Xiangxin Meng

Additional Contributors
Xinchen Wang
Ruida Hu
Yuanan Xiao
Yizhou Liu
Zhao Zhang

Special Thanks
Junjie Chen
Cuiyun Gao
Yun Lin
Yingfei Xiong

Project Lead
Chao Peng

Team Management
Xia Liu

[*] Equal contribution.

References

I. Abdelaziz, K. Basu, M. Agarwal, S. Kumaravel, M. Stallone, R. Panda, Y. Rizk, G. Bhargav, M. Crouse, C. Gunasekara, et al. Granite-function calling model: Introducing function calling abilities via multi-task learning of granular tasks. *arXiv preprint arXiv:2407.00121*, 2024.

G. AI. Gemini developer api. https://ai.google.dev/gemini-api/docs/, 2025.

Anthropic. Build with claude. https://www.anthropic.com/api, 2025a.

Anthropic. Claude 3.7 sonnet and claude code. https://www.anthropic.com/news/claude-3-7-sonnet, 2025b.

S. bench Team. Swe-bench leaderboard. https://www.swebench.com/, 2025.

M. Bordese. Unified diff python parsing/metadata extraction library. https://github.com/matiasb/python-unidiff, 2025.

I. Bouzenia, P. Devanbu, and M. Pradel. RepairAgent: An Autonomous, LLM-Based Agent for Program Repair . In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*, pages 2188–2200. IEEE, 2025.

L. Breiman. Bagging predictors. *Machine learning*, 24:123–140, 1996.

R. P. Buse and W. R. Weimer. Learning a metric for code readability. *IEEE Transactions on software engineering*, 36(4):546–558, 2009.

T. Chen and C. Flaherty. #1 open-source agent on swe-bench verified by combining claude 3.7 and o1. https://www.augmentcode.com/blog/1-open-source-agent-on-swe-bench-verified-by-combining-claude-3-7-and-o1, 2025.

Y. Chen, Y. Wang, Y. Gu, X. He, P. Ghamisi, and X. Jia. Deep learning ensemble for hyperspectral image classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(6): 1882–1897, 2019.

J. Clarke, J. J. Dolado, M. Harman, R. Hierons, B. Jones, M. Lumkin, B. Mitchell, S. Mancoridis, K. Rees, M. Roper, et al. Reformulating software engineering as a search problem. *IEE Proceedings-software*, 150(3): 161–175, 2003.

A. Code. Augment swe-bench verified agent. https://github.com/augmentcode/augment-swebench-agent, 2025.

I. Cohen, Y. Huang, J. Chen, J. Benesty, J. Benesty, J. Chen, Y. Huang, and I. Cohen. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4, 2009.

CrossHair. An analysis tool for python that blurs the line between testing and type systems. https://github.com/pschanely/CrossHair, 2025.

G. DeepMind. Gemini 2.5 pro. https://deepmind.google/models/gemini/pro/, 2025.

G. A. Di Lucca and M. Di Penta. Integrating static and dynamic analysis to improve the comprehension of existing web applications. In *Seventh IEEE International Symposium on Web Site Evolution*, pages 87–94. IEEE, 2005.

S. Džeroski and B. Ženko. Is combining classifiers with stacking better than selecting the best one? *Machine learning*, 54:255–273, 2004.

T. Eisenbarth, R. Koschke, and D. Simon. Aiding program comprehension by static and dynamic feature analysis. In *Proceedings IEEE International Conference on Software Maintenance. ICSM 2001*, pages 602–611. IEEE, 2001.

S. Fakhoury, S. Chakraborty, M. Musuvathi, and S. K. Lahiri. Towards generating functionally correct code edits from natural language issue descriptions. *arXiv preprint arXiv:2304.03816*, 2023.

Z. Fan, X. Gao, M. Mirchev, A. Roychoudhury, and S. H. Tan. Automated repair of programs from large language models. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 1469–1481. IEEE, 2023.

Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.

L. Guo, W. Tao, R. Jiang, Y. Wang, J. Chen, X. Liu, Y. Ma, M. Mao, H. Zhang, and Z. Zheng. Omnigirl: A multilingual and multimodal benchmark for github issue resolution. *Proceedings of the ACM on Software Engineering*, 2(ISSTA):24–46, 2025.

X. Han, X. Chen, and L.-P. Liu. Gan ensemble for anomaly detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4090–4097, 2021.

M. Harman and B. F. Jones. Search-based software engineering. *Information and software Technology*, 43(14): 833–839, 2001.

C. E. Jimenez, J. Yang, A. Wettig, S. Yao, K. Pei, O. Press, and K. R. Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *The Twelfth International Conference on Learning Representations*, 2023.

M. G. Kendall. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93, 1938.

J. Kim, J. Jeon, S. Hong, and S. Yoo. Predictive mutation analysis via the natural language channel in source code. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(4):1–27, 2022.

M. Kintis, M. Papadakis, Y. Jia, N. Malevris, Y. Le Traon, and M. Harman. Detecting trivial mutant equivalences via compiler optimisations. *IEEE Transactions on Software Engineering*, 44(4):308–333, 2017.

D. Li, S. Cao, C. Cao, X. Li, S. Tan, K. Keutzer, J. Xing, J. E. Gonzalez, and I. Stoica. S*: Test time scaling for code generation. *arXiv preprint arXiv:2502.14382*, 2025.

X. Li and L. Zhang. Transforming programs and tests in tandem for fault localization. *Proceedings of the ACM on Programming Languages*, 1(OOPSLA):1–30, 2017.

F. Lin, D. J. Kim, and T.-H. P. Chen. Soen-101: Code generation by emulating software process models using large language model agents. In *2025 IEEE/ACM 46th International Conference on Software Engineering (ICSE)*, 2025.

J. Liu, C. S. Xia, Y. Wang, and L. Zhang. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems*, 36: 21558–21572, 2023.

W. Liu, X. Huang, X. Zeng, xinlong hao, S. Yu, D. Li, S. Wang, W. Gan, Z. Liu, Y. Yu, Z. WANG, Y. Wang, W. Ning, Y. Hou, B. Wang, C. Wu, W. Xinzhi, Y. Liu, Y. Wang, D. Tang, D. Tu, L. Shang, X. Jiang, R. Tang, D. Lian, Q. Liu, and E. Chen. ToolACE: Winning the points of LLM function calling. In *The Thirteenth International Conference on Learning Representations*, 2025.

Y. Liu, P. Gao, X. Wang, J. Liu, Y. Shi, Z. Zhang, and C. Peng. Marscode agent: Ai-native automated bug fixing. *arXiv preprint arXiv:2409.00899*, 2024.

W. Ma, S. Liu, Z. Lin, W. Wang, Q. Hu, Y. Liu, C. Zhang, L. Nie, L. Li, and Y. Liu. Lms: Understanding code syntax and semantics for code analysis. *arXiv preprint arXiv:2305.12138*, 2023.

T. Mahmud, B. Duan, C. Pasareanu, and G. Yang. Enhancing llm code generation with ensembles: A similarity-based selection approach. *arXiv preprint arXiv:2503.15838*, 2025.

OpenAI. Introducing swe-bench verified. https://openai.com/index/introducing-swe-bench-verified/, 2024.

OpenAI. Introducing gpt-4.1 in the api. https://openai.com/index/gpt-4-1/, 2025a.

OpenAI. The fastest and most powerful platform for building ai products. https://openai.com/api/, 2025b.

M. Papadakis, Y. Jia, M. Harman, and Y. Le Traon. Trivial compiler equivalence: A large scale empirical study of a simple, fast and effective equivalent mutant detection technique. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, volume 1, pages 936–946. IEEE, 2015.

R. Polikar. Ensemble learning. *Ensemble machine learning: Methods and applications*, pages 1–34, 2012.

S. Ren, D. Guo, S. Lu, L. Zhou, S. Liu, D. Tang, N. Sundaresan, M. Zhou, A. Blanco, and S. Ma. Codebleu: a method for automatic evaluation of code synthesis. *arXiv preprint arXiv:2009.10297*, 2020.

M. Renze. The effect of sampling temperature on problem solving in large language models. In *Findings of the association for computational linguistics: EMNLP 2024*, pages 7346–7356, 2024.

H. Ruan, Y. Zhang, and A. Roychoudhury. Specrover: Code intent extraction via llms. In *2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE)*, pages 963–974. IEEE, 2025.

O. Sagi and L. Rokach. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 8(4):e1249, 2018.

B. Sebastian, C. Christian, and P. Alexander. Predicting the resilience of obfuscated code against symbolic execution attacks via machine learning. In *Proceedings of the 26th USENIX Security Symposium (USENIX Security 17), Vancouver, BC, Canada*, pages 16–18, 2017.

D. Shin, S. Yoo, M. Papadakis, and D.-H. Bae. Empirical evaluation of mutation-based test case prioritization techniques. *Software Testing, Verification and Reliability*, 29(1-2):e1695, 2019.

C. V. Snell, J. Lee, K. Xu, and A. Kumar. Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth International Conference on Learning Representations*, 2025.

C. Spearman. The proof and measurement of association between two things. 1961.

Z. Tian, H. Shu, D. Wang, X. Cao, Y. Kamei, and J. Chen. Large language models for equivalent mutant detection: How far are we? In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 1733–1745, 2024.

Z. Tian, J. Chen, and X. Zhang. Fixing large language models' specification misunderstanding for better code generation. In *2025 IEEE/ACM 46th International Conference on Software Engineering (ICSE)*, 2025.

T. Titcheu Chekam, M. Papadakis, T. F. Bissyandé, Y. Le Traon, and K. Sen. Selecting fault revealing mutants. *Empirical Software Engineering*, 25:434–487, 2020.

Trae-Agent. `https://github.com/bytedance/trae-agent`, 2025.

R. Wang, J. Guo, C. Gao, G. Fan, C. Y. Chong, and X. Xia. Can llms replace human evaluators? an empirical study of llm-as-a-judge in software engineering. *Proceedings of the ACM on Software Engineering*, 2(ISSTA): 1955–1977, 2025a.

X. Wang, B. Li, Y. Song, F. F. Xu, X. Tang, M. Zhuge, J. Pan, Y. Song, B. Li, J. Singh, H. H. Tran, F. Li, R. Ma, M. Zheng, B. Qian, Y. Shao, N. Muennighoff, Y. Zhang, B. Hui, J. Lin, R. Brennan, H. Peng, H. Ji, and G. Neubig. Openhands: An open platform for AI software developers as generalist agents. In *The Thirteenth International Conference on Learning Representations*, 2025b.

Y. Wang, M. Pradel, and Z. Liu. Are" solved issues" in swe-bench really solved correctly? an empirical study. *arXiv preprint arXiv:2503.15223*, 2025c.

F. Wilcoxon, S. Katti, R. A. Wilcox, et al. *Critical values and probability levels for the Wilcoxon rank sum test and the Wilcoxon signed rank test*, volume 1. American Cyanamid Pearl River, NY, 1963.

C. S. Xia, Y. Deng, S. Dunn, and L. Zhang. Demystifying llm-based software engineering agents. *Proc. ACM Softw. Eng.*, 2(FSE), 2025.

J. Yang, C. E. Jimenez, A. Wettig, K. Lieret, S. Yao, K. Narasimhan, and O. Press. Swe-agent: Agent-computer interfaces enable automated software engineering. *Advances in Neural Information Processing Systems*, 37: 50528–50652, 2024.

H. Zhang and M. O. Shafiq. Survey of transformers and towards ensemble learning using transformers for natural language processing. *Journal of big Data*, 11(1):25, 2024.

K. Zhang, W. Yao, Z. Liu, Y. Feng, Z. Liu, R. RN, T. Lan, L. Li, R. Lou, J. Xu, et al. Diversity empowers intelligence: Integrating expertise of software engineering agents. In *The Thirteenth International Conference on Learning Representations*, 2024a.

K. Zhang, H. Zhang, G. Li, J. You, J. Li, Y. Zhao, and Z. Jin. Sealign: Alignment training for software engineering agent. *arXiv preprint arXiv:2503.18455*, 2025.

Y. Zhang and A. Mesbah. Assertions are strongly correlated with test suite effectiveness. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, pages 214–224, 2015.

Y. Zhang, H. Ruan, Z. Fan, and A. Roychoudhury. Autocoderover: Autonomous program improvement. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 1592–1604, 2024b.

Y. Zhu, J. Li, G. Li, Y. Zhao, Z. Jin, and H. Mei. Hot or cold? adaptive temperature sampling for code generation with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 437–445, 2024.

A. Örwall. Moatless tools. `https://github.com/aorwall/moatless-tools`, 2025.