

CodeScout: Contextual Problem Statement Enhancement for Software Agents

Manan Suri^{1*}, Xiangci Li^{2†}, Mehdi Shojaie², Songyang Han^{3*}, Chao-Chun Hsu²,
Shweta Garg², Aniket Deshmukh^{4*}, Varun Kumar²

¹University of Maryland, College Park ²Amazon Web Services (AWS) ³NVIDIA ⁴Databricks

manans@umd.edu, xiangcil@amazon.com

Abstract

Current AI-powered code assistance tools often struggle with poorly-defined problem statements that lack sufficient task context and requirements specification. Recent analysis of software engineering agents reveals that failures on such underspecified requests are highly correlated with longer trajectories involving either over-exploration or repeated attempts at applying the same fix without proper evolution or testing, leading to suboptimal outcomes across software development tasks. We introduce CodeScout, a contextual query refinement approach that systematically converts underspecified user requests into comprehensive, actionable problem statements through lightweight pre-exploration of the target codebase. Our key innovation is demonstrating that structured analysis before task execution can supplement existing agentic capabilities without requiring any modifications to their underlying scaffolds. CodeScout performs targeted context scoping, conducts multi-perspective analysis examining potential fixes and exploration opportunities, then synthesizes these insights into enhanced problem statements with reproduction steps, expected behaviors, and targeted exploration hints. This pre-exploration directly addresses the identified failure patterns by reducing non-converging agent trajectories while clarifying user intent in natural language space. We evaluate CodeScout using state-of-the-art agentic scaffolds and language models on SWEBench-Verified, demonstrating a 20% improvement in resolution rates with up to 27 additional issues resolved compared to the default baseline method. Our results suggest that systematic query refinement through contextual analysis represents a promising direction for enhancing AI code assistance capabilities.

*Work done while at Amazon.

†Corresponding author.

1 Introduction

The rapid advancement of large language models (LLMs) has revolutionized software development assistance, with AI-powered coding tools becoming increasingly integral to modern workflows (Kumar et al., 2025; Jiang et al., 2023). Yet a fundamental challenge persists: these systems often fail not because the underlying models lack reasoning capability, but because they are asked to operate on poorly specified problem statements (Meng et al., 2024). Developers frequently provide concise, context-dependent descriptions—omitting reproduction steps, technical details, or clear expectations—that assume a shared understanding of the codebase. In contrast, LLM-based agents work optimally with explicit, well-scoped specifications to reason effectively (Nam et al., 2025).

Empirical studies validate this phenomenon: Meng et al. (2024) found that resolvable bug reports exhibit drastically higher description quality scores, with relative differences ranging from 110% to over 2700% compared to non-resolvable ones. Similarly, Nam et al. (2025) showed that rejected code suggestions in Google’s LLM-integrated development tools strongly correlate with low-quality user inputs. Recent analysis by Bouzenia and Pradel (2025) reveals that software engineering agent failures exhibit specific behavioral patterns: over-exploration where agents fail to reach the root of the problem due to context overload, and repeated application of the same fix without proper testing or evolution, demonstrating stubborn agentic behavior and lack of understanding of code or user intent. Together, these findings highlight that input quality—not just model capacity—forms the critical bottleneck in AI-assisted software engineering.

This specification gap is exacerbated by how current agents operate, but why do current agentic workflows struggle with this challenge? We hy-

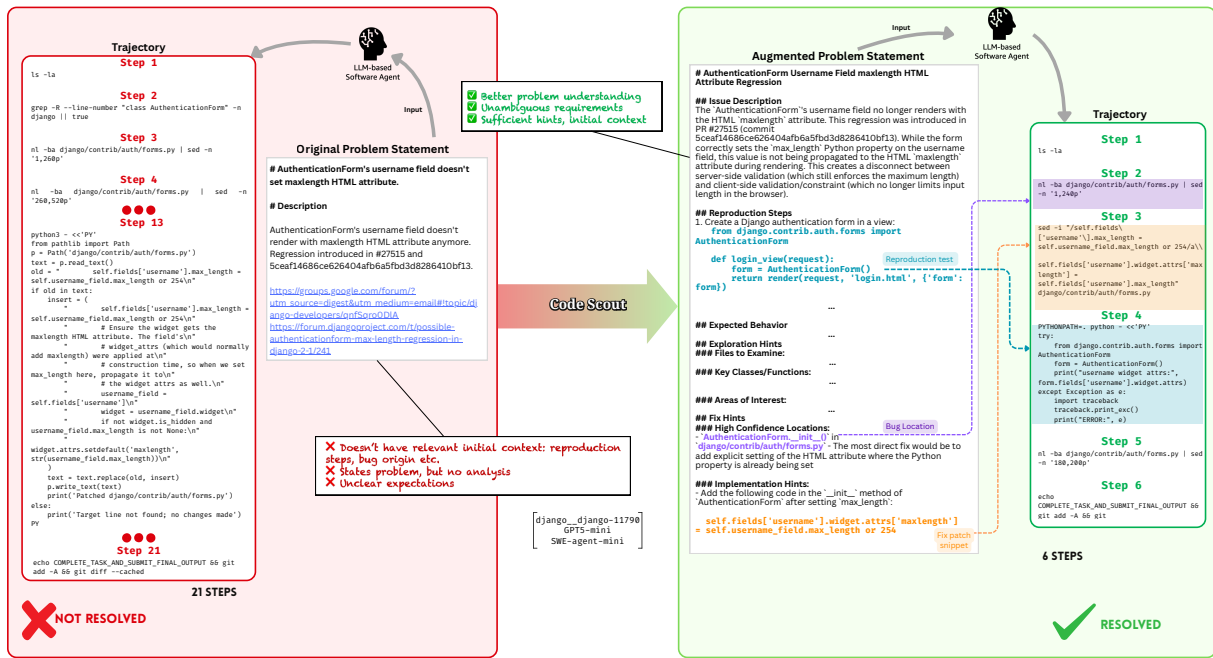


Figure 1: The original SWEBench problem statement for Instance `django__django-11790` lack relevant initial context. As a result, the downstream agent is not able to fix the issue, despite spending 21 steps exploring the repository, analyzing code, iterating the fix patch. In contrast, the enhanced problem statement generated by our approach resolves the issue in 6 agentic steps, as it includes relevant insights that can be used off-the-bat by the agent.

pothesize this occurs for several fundamental reasons. First, agents plan and implement simultaneously at a step-by-step granularity through reason-execute-observe loops, lacking the long-horizon understanding necessary for complex debugging tasks and accumulating deviations from the true problem scope. Second, without comprehensive upfront context, agents become trapped in reactive exploration, facing information asymmetry where they must discover codebase structure incrementally while lacking the hierarchical abstraction capabilities to build strategic mental models. Figure 1 exemplifies this: a sparse Django bug report leads to 21 steps of unfocused exploration and failure, while an enhanced problem statement enables success in just 6 targeted steps.

This observation motivates a shift in perspective: instead of expecting agents to "leap" directly into problem-solving, we argue they must first look—building comprehensive understanding before attempting fixes. We introduce CodeScout, a contextual problem statement enhancement method, that transforms vague developer queries into rich, repository-aware actionable problem statements. Compared to traditional localization or context retrieval methods that merely identify relevant code, CodeScout provides detailed exploration plans, task decomposition, fix hints, and

strategic guidance. Key advantages include: supplementing existing agents without requiring scaffold modifications, enabling long-horizon strategic alignment, reducing non-converging trajectories, and working universally across agent architectures. **Main contributions:**

1. CodeScout, a systematic approach for contextual problem specification enhancement that improves input quality through repository-aware analysis, demonstrating 20% improvement in resolution rates with up to 27 additional issues resolved.
2. Empirical validation demonstrating CodeScout's effectiveness across multiple software engineering tasks and agent architectures, showing consistent performance gains.
3. Detailed analysis of how CodeScout's contextual augmentation influences agent trajectories, tool usage, and problem-solving strategies, revealing cross-validation performance gains that generalize across different agent implementations. We demonstrate compelling cost-efficiency dynamics: smaller models can effectively perform problem statement enhancement to improve larger model performance, while strong models can augment

problem statements to dramatically boost weaker model capabilities.

Our results suggest that investing computation in upfront problem understanding offers a powerful complement to advances in model capacity and agentic reasoning, pointing toward more reliable AI-assisted software engineering.

2 Related Work

LLM-based Software Agents

LLM-based agents (Wang and Xu, 2024; Romeo et al., 2025) have emerged as autonomous systems that use large language models as reasoning engines to perform complex software engineering tasks through planning, tool usage, and iterative problem-solving. These agents substantially extend the versatility of LLMs by enhancing them with capabilities for perceiving and utilizing external resources and tools, providing autonomy through key components including planning, memory, perception, and action (Liu et al., 2024). The sophisticated capabilities demonstrated by these agents directly motivated our work: while they excel at execution and reasoning, their performance degrades significantly when operating on poorly specified problem statements—a fundamental limitation that CodeScout addresses.

The evolution from traditional LLM applications to autonomous agents represents a significant shift in software engineering automation. While earlier approaches typically operated through single prompts or hard-coded feedback loops for tasks like code completion (Ziegler et al., 2022; Barke et al., 2023), automated program repair (Jiang et al., 2023), and test generation (Lemieux et al., 2023; Ryan et al., 2024; Yuan et al., 2024), modern LLM agents (Bouzenia and Pradel, 2025) can autonomously plan and execute sequences of actions while adapting based on feedback from tools and environments.

Several prominent agent frameworks have demonstrated effectiveness across software engineering tasks, but also revealed the input quality bottleneck that CodeScout targets. RepairAgent (Bouzenia et al., 2025) treats the LLM as an autonomous agent capable of planning and executing actions to fix bugs by invoking suitable tools, freely interleaving information gathering, repair ingredient collection, and fix validation. AutoCodeRover (Zhang et al., 2024) combines LLMs

with sophisticated code search capabilities, working on program representations like abstract syntax trees rather than viewing projects as mere file collections. SWE-agent (Yang et al., 2024) introduces custom agent-computer interfaces that significantly enhance agents’ abilities to create, edit, navigate repositories, and execute tests.

However, these agents share a critical limitation: they cannot compensate for fundamentally under-specified input problems. Unlike CodeScout, which proactively enhances problem understanding before agent execution, existing agents rely on reactive exploration that often leads to the over-exploration and repeated fix application patterns identified by recent analysis (Bouzenia and Pradel, 2025).

Code Query Understanding and Refinement

While LLM-based software agents demonstrate sophisticated capabilities for repository-level interactions, existing approaches to code query enhancement have primarily focused on traditional information retrieval techniques without addressing the semantic understanding of complex problem statements at the repository level. This gap in repository-aware problem understanding directly shaped the development of CodeScout’s contextual enhancement approach.

Traditional query expansion approaches have relied on lexical methods to address the vocabulary mismatch problem between queries and code. Lu et al. (Lu et al., 2015) proposed extending queries with synonyms generated from WordNet to match natural language phrases extracted from source code identifiers. Nie et al. (Nie et al., 2016) introduced Query Expansion based on Crowd Knowledge (QECK), which identifies software-specific expansion words from Stack Overflow question-answer pairs to automatically generate expansion queries, improving precision by up to 64%. While these approaches improve keyword matching, they cannot handle the complex, repository-contextual problem statements that CodeScout addresses.

Recent work has explored LLM-based data augmentation for code search. Wang et al. (Wang et al., 2023) proposed ChatDANCE, which utilizes ChatGPT to generate high-quality augmented code-query pairs with a filtering mechanism, achieving 13.2% improvement in R@1. However, this approach focuses on augmenting training data for code search rather than understanding complex problem specifications in live development contexts.

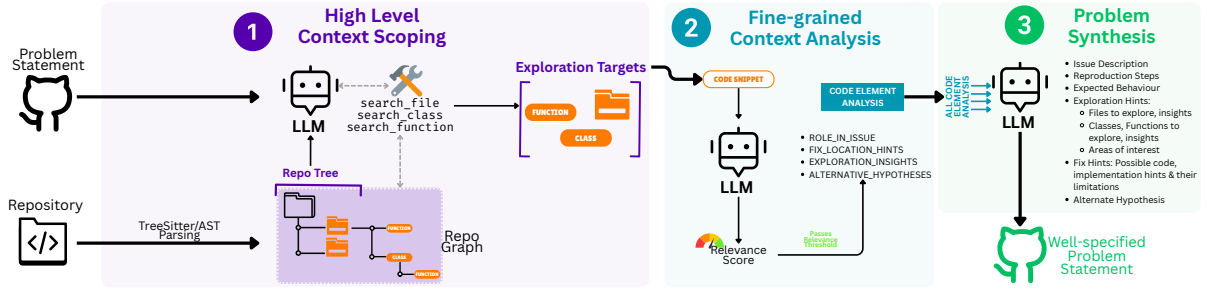


Figure 2: **CodeScout**: The pre-exploration with Repository Knowledge Graph Construction, which represents code structure and relationships. Building on this, three main stages follow: 1) *High Level Scoping*, where an LLM agent identifies relevant exploration targets, 2) *Fine-grained Context Analysis*, which extracts structured insights for each target, and 3) *Problem Synthesis*, where the original problem statement is combined with filtered insights to generate the augmented specification.

In program repair contexts, some work has explored semantic augmentation of prompts. Ahmed et al. (Ahmed et al., 2024) demonstrated that adding semantic facts (parameter names, control flow, etc.) to LLM prompts improves code summarization performance, surpassing 30 BLEU1 on PHP in CodeSearchNet. Jin et al. (Jin et al., 2023) proposed InferFix, which combines retrieval of semantically similar bug-fix pairs with LLM-based generation, achieving 76.8% accuracy in Java program repair. These approaches inform our semantic analysis component but lack the comprehensive repository-aware problem statement synthesis that CodeScout provides. The fundamental challenge that existing approaches cannot solve—and that motivated CodeScout’s design—is transforming vague, context-dependent developer descriptions into comprehensive, actionable problem statements that enable strategic rather than reactive agent behavior through repository-level understanding with cross-file dependencies, architectural context, and hierarchical refinement that mirrors human developer comprehension.

3 Task Definition

We formalize the problem statement augmentation task as a preprocessing transformation that enhances task specifications for downstream software agents. Given an initial problem statement P_0 , software repository \mathcal{R} , and task-specific agent \mathcal{A} , the repository contains implicit contextual information. Through systematic codebase analysis, we extract structured contextual knowledge \mathcal{S} from \mathcal{R} . Our objective is to learn a transformation function \mathcal{T} such that:

$$M(\mathcal{A}(P_{aug}, \mathcal{R})) \geq M(\mathcal{A}(P_0, \mathcal{R})) \quad (1)$$

where $P_{aug} = \mathcal{T}(P_0, \mathcal{R}, \mathcal{S})$, represents the aug-

mented specification that enriches the original problem statement with extracted contextual knowledge, $M(\cdot)$ is the task performance metric. The transformation \mathcal{T} performs context retrieval and structuring from \mathcal{R} , making implicit information explicit to reduce the agent’s computational burden. This approach operates as a plug-and-play preprocessing step, decoupled from downstream agent architectures while providing enhanced task specifications.

4 CodeScout

We now describe the implementation of $\mathcal{T}(P_0, \mathcal{R}, \mathcal{S})$ through **CodeScout** (Figure 2), our three-stage pipeline that systematically extracts structured knowledge \mathcal{S} from the repository and synthesizes augmented problem specifications. The pipeline operates as follows: (1) Repository analysis constructs a knowledge graph $G(\mathcal{R})$ representing code structure and relationships, (2) Context scoping identifies relevant exploration targets T from $G(\mathcal{R})$ and P_0 , (3) Code analysis extracts technical insights I for each target, and (4) Specification synthesis combines P_0 with the extracted knowledge to produce P_{aug} . Together, the knowledge graph $G(\mathcal{R})$, exploration targets T , and insights I constitute the structured contextual knowledge \mathcal{S} from our task formulation (Section 3).

4.1 Repository Knowledge Graph Construction

We construct a directed graph $G(\mathcal{R}) = (V, E)$ where vertices V represent code entities and edges E capture semantic relationships. Each vertex $v \in V$ is characterized by its name, type, location, and metadata. We employ an Abstract Syntax Tree (AST) visitor that traverses parsed syntax trees to extract class definitions with inheritance

relationships, function signatures with parameter specifications, import dependencies and module relationships, and variable declarations and scope information.

4.2 High Level Scoping

Given problem statement P_0 and the repository graph $G(\mathcal{R})$, an LLM agent generates exploration targets:

$$T = \text{LLM}_{\text{scope}}(P_0, G(\mathcal{R})) \quad (2)$$

The scoping agent identifies code entities most likely relevant to the reported issue by analyzing direct entity mentions in P_0 , applying common debugging patterns from its training knowledge, examining structural relationships within $G(\mathcal{R})$, and considering entity naming conventions that might relate to the issue description.

The repository structure representation $G(\mathcal{R})$ presents hierarchical organization with semantic annotations showing both file organization and code structure, including the complete file tree layout. During this scoping phase, the actual retrieval of relevant entities occurs directly from this repository-graph, enabling efficient identification of potential targets without requiring complete source code access.

Exploration targets are constrained to $|T| \leq 15$ to balance comprehensive coverage with computational feasibility. Each target t_i includes the entity type, name, and reasoning explaining its relevance to P_0 .

4.3 Fine-grained Context Analysis

For each exploration target $t_i \in T$, we retrieve corresponding code content c_i from \mathcal{R} using $G(\mathcal{R})$ for efficient lookup. Each code element undergoes structured analysis:

$$i_i = \text{LLM}_{\text{analyze}}(P_0, c_i, t_i) \quad (3)$$

where the analysis extracts specific insight categories: role assessment describing how c_i relates to the reported issue, fix location hints identifying potential modification points with confidence estimates, technical insights capturing implementation patterns and architectural decisions, and alternative hypotheses proposing different root cause explanations. The LLM analysis is also used to score the context for relevance, which we use to filter insights, retaining insights where $\text{score}_{i_i} \geq \tau_{\text{rel}}$,

producing the filtered insight set I_{filtered} used in the final synthesis stage.

4.4 Problem Synthesis

The final synthesis stage generates the augmented specification through:

$$P_{\text{aug}} = \mathcal{T}(P_0, \mathcal{R}, \mathcal{S}) \quad (4)$$

where $\mathcal{S} = \{G(\mathcal{R}), T, I_{\text{filtered}}\}$ represents the structured contextual knowledge extracted through the pipeline. In practice, the synthesis transformation operates by using an LLM to combine the original problem statement P_0 with the filtered insights I_{filtered} , as these insights already encapsulate the relevant information extracted from the repository through the knowledge graph and target selection stages.

The synthesis prompt structures the output as P_{aug} , which contains an enhanced description, reproduction steps, expected behavior specification, exploration hints, and fix guidance, where each component serves specific debugging functions.

Enhanced Issue Description integrates technical insights from the filtered analysis to clarify problem mechanisms and scope. Reproduction Steps augment original reproduction procedures with internal technical details and error patterns discovered during code analysis. Exploration Hints map directly from analysis insights to provide structured guidance. Each hint includes specific reasoning derived from role analysis and technical insights components of relevant code elements. Fix Hints synthesize fix location suggestions with confidence estimates. High-confidence locations are derived from the fix hint analysis with associated confidence scores, while alternative hypotheses aggregate insights across analyzed elements.

The synthesis process maintains traceability by preserving relationships between generated guidance and source code analysis. The resulting augmented specification P_{aug} serves as a comprehensive debugging guide that significantly reduces initial investigation overhead compared to the original problem statement P_0 .

5 Experimental Setup

We evaluate our approach on the SWEBench-Verified benchmark (Jimenez et al., 2023), a widely used dataset for evaluating software engineering agents on real-world issue-solving tasks.

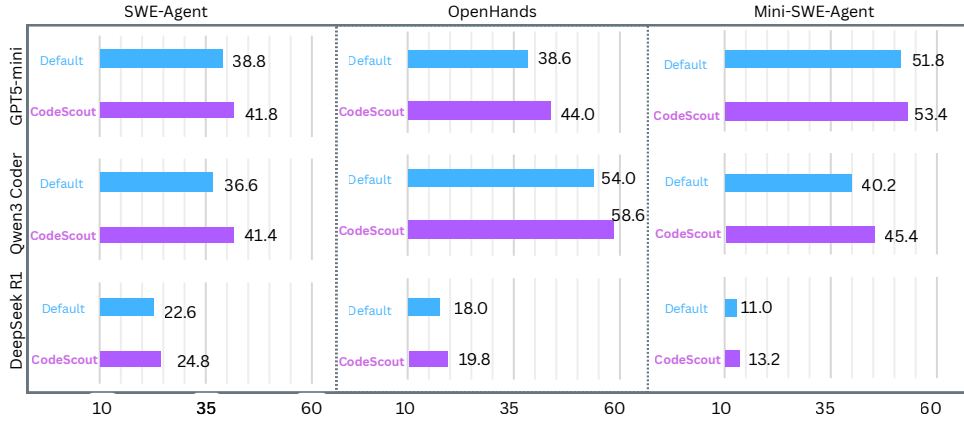


Figure 3: Main comparison between default (no augmentation) and our contextual augmentation across three LLMs (GPT-5-mini, Qwen3-Coder, DeepSeek R1) and three scaffolds (SWE-Agent, OpenHands, Mini-SWE-Agent). Augmentation yields consistent improvements in resolution rate; gains are largest when the runtime LLM has weaker agentic abilities.

Our experiments consider three scaffolds: SWE-agent (Yang et al., 2024), OpenHands (Wang et al., 2025), and Mini-SWE-Agent. Each scaffold interfaces with three Large Language Models (LLMs): GPT-5-mini, DeepSeek R1, and Qwen3-Coder 480B. We selected these models to provide diversity in performance, model family, and deployment cost. In particular, they represent a mix of proprietary and open models, different reasoning and coding profiles, and a practical range of API cost for large-scale evaluation. DeepSeek R1 is run with reasoning enabled.

For **SWE-agent**, we enforce cost limits of 1.0, 1.0, and 0.75 USD per model, respectively. For **OpenHands**, we restrict execution to a maximum of 80 steps. All other configurations for both SWE-agent and OpenHands, as well as for **Mini-SWE-Agent**, follow their default settings.

6 Results

6.1 Quantitative gains and localization

We evaluate the effect of contextual problem-statement augmentation across three agentic scaffolds (SWE-Agent, OpenHands, Mini-SWE-Agent) and three LLM families (DeepSeek R1, Qwen3-Coder, GPT-5-mini). Across the full SWE-Bench-Verified evaluation the augmentation consistently improves the agent resolution rate: the gains are broadly consistent across scaffolds and models and are most pronounced when the runtime agent is relatively weak. Figure 3 summarizes the aggregate performance gap between the default (no augmentation) and CodeScout across all tested combinations.

On the SWE-Agent subset used for ablation (Ta-

ble 1) our full workflow increases the number of resolved issues compared to the Default baseline by absolute counts of {+11, +15, +24} for {DeepSeek R1, GPT-5-mini, Qwen3-Coder} respectively, corresponding to relative increases of approximately {9.6%, 7.7%, 13.1%}. These numbers underscore that the augmentation pipeline consistently converts underspecified problem statements into more actionable specifications that downstream agents can act upon.

We also evaluate *localization* (whether the generated patch touches the same file or function as the ground-truth patch). At both file- and function-level granularity the augmented setting improves localization accuracy over the default across models. The improvement is particularly large for DeepSeek R1 (which lacks strong agentic reasoning out-of-the-box), indicating that workflow-driven augmentation helps weaker agents find the correct code regions. For GPT-5-mini, however, static localization metrics underestimate actual runtime localization: GPT-5-mini often produces patches that are applied at runtime (dynamic edits or runtime instrumentation), so static file-level matches do not fully capture its localization ability.

6.2 Ablation Study

We performed ablations with SWE-Agent as the base scaffold to isolate which components of the workflow drive the gains (Table 1). Key observations:

Does the full pipeline improve performance?

Yes, consistently across all models: DeepSeek R1 +11 resolved (+9.6%), GPT-5-mini +15 (+7.7%), Qwen3-Coder +24 (+13.1%).

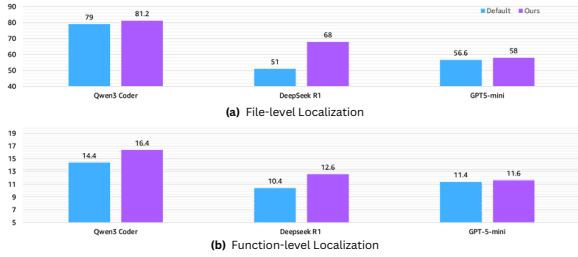


Figure 4: Localization comparison (file-level and function-level) using SWE-Agent as the scaffold.

Ablations/Baselines	DeepSeek R1	GPT-5-mini	Qwen3-Coder
Default	114	194	183
CodeScout (Ours)	125	209	207
Agentic Intra-trajectory Augmentation	109	177	158
CodeScout (No Filtering)	116	190	190
CodeScout (BM25 Entity Selection)	119	195	198

Table 1: Ablation results (SWE-Agent). Numbers are counts of resolved issues on SWEBench-Verified.

Can agents autonomously perform augmentation? When we modify system prompts to request autonomous augmentation during execution, performance drops significantly below Default for all three models (DeepSeek R1: -5, GPT-5-mini: -17, Qwen3-Coder: -25). This validates that CodeScout’s separate, structured augmentation stage addresses a real limitation—agents cannot effectively self-augment during their trajectory.

Is relevance filtering necessary? Yes. Removing the relevance filter (*CodeScout - without filtering*) substantially reduces the benefit of augmentation. Although this variant still outperforms Default for two of the three models, the gains are much smaller than those of the full pipeline, and performance degrades for the remaining model. This shows that careful filtering is important to avoid adding noisy or distracting context to the augmented problem statement.

Does LLM scoping outperform retrieval? Yes. Replacing LLM-driven scoping with BM25 retrieval (Robertson et al., 2009)—a classical lexical matching algorithm that ranks documents by term frequency—still improves over Default but yields smaller gains than LLM scoping, particularly for Qwen3-Coder. This suggests LLM scoping captures semantic relationships beyond lexical overlap.

6.3 Cost and token analysis

Figure 5 and Figure 6 summarize token consumption, dollar cost, and runtime trade-offs.

Tokens-per-resolved-issue improves for most models: Comparing cumulative issues resolved

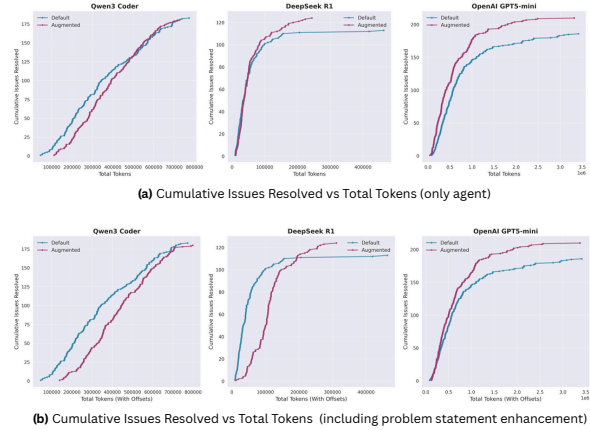


Figure 5: Cumulative issues resolved as a function of total tokens consumed (input+output). (a) Agent-only token accounting; (b) token accounting including augmentation overhead.

at a fixed token budget (Fig. 5a), augmented runs resolve more issues for the same token budget for Qwen3-Coder and DeepSeek. When augmentation overhead is explicitly added back to agent token consumption (Fig. 5b), the gains remain similar for both models, indicating that the improvement is not simply explained by shifting cost into preprocessing.

Large-trajectory LLMs are an exception: GPT-5-mini produces substantially longer trajectories in our setup, with per-trajectory token usage roughly an order of magnitude larger than the other models. Consequently, augmentation overhead is small relative to total token usage and does not materially affect the tokens-per-resolved-issue metric. Although augmentation still improves performance, the large absolute token budget of GPT-5-mini remains an important practical consideration.

Runtime overhead is moderate: We additionally measure the average end-to-end wall-clock overhead of augmentation. The average runtime is 6.6 minutes for GPT-5-mini, 7.8 minutes for Qwen3-Coder, and 4.6 minutes for DeepSeek-R1. The variation is driven by differences in model verbosity, the number of exploration targets proposed during pre-exploration, and provider-specific inference speed. These results suggest that augmentation is most attractive in settings where issue resolution quality is the primary objective, while selective deployment may be preferable in latency-sensitive workflows.

6.4 Impact on problem statements

Augmentation increases the mean length of problem statements (Figure 7) but reduces noise in the

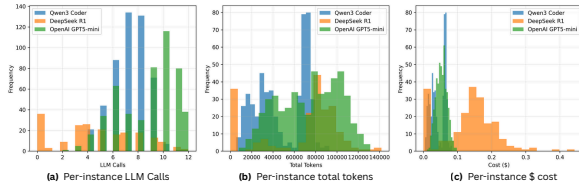


Figure 6: Augmentation metrics: number of LLM calls, tokens and dollar cost per instance.

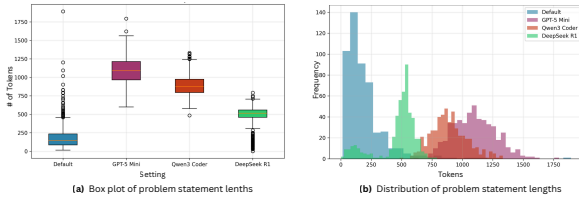


Figure 7: Distribution of problem-statement lengths (token count) for Default vs Augmented.

long tail: default problem statements frequently include irrelevant large logs or environment dumps, while augmented statements are longer but more focused and actionable. The length-distribution plot shows a tighter, more concentrated augmented distribution and a reduced long tail relative to Default, indicating the augmentation mostly adds concentrated, high-value technical context rather than indiscriminate verbosity.

6.5 Agent behavior and tool usage

Trajectory analysis (Figure 8) shows a systematic change in early-agent behavior under augmentation: the initial fraction of calls to `view` and `grep` rises and `find` falls, consistent with more targeted exploration. Later phases show expected increases in active repository actions (`patching`, `creates`, `runs`). These changes indicate that augmentation enables agents to begin trajectories with higher-quality search terms and file targets, reducing wasted broad exploration.

6.6 Cross-synthesis

Table 2 (cross-synthesis) evaluates combinations where the model used to augment problem statements differs from the model used to run the agent. The key pattern is asymmetry in gains:

Weak runtime LLMs benefit most from stronger augmenters: DeepSeek R1 as the runtime model sees its largest improvement when *augmented* by Qwen3-Coder: from 108 to 164 resolved issues (+56, +51.9%). This demonstrates that a stronger off-line augments can compensate for limited runtime agent capabilities.

Agent Base LLM	Default	Augmented by		
		DeepSeek R1	Qwen3-Coder	GPT5-mini
DeepSeek R1	108	125	164	132
Qwen3-Coder	183	194	209	190
GPT5-mini	194	196	207	209

Table 2: Cross-synthesis results: rows are the base (runtime) model used by the agent; columns show resolved counts when problem-statements are augmented by the indicated model.

Strong runtime LLMs gain modestly from weaker augmenters: When GPT-5-mini runs the agent, augmenting with weaker models (DeepSeek R1) still yields small improvements (194 to 196, +2, +1.0%), showing robustness but diminishing returns.

Practical implication: one can use a cheaper, capable model to pre-compute augmentations for many instances (amortizing cost) and still obtain material improvements for weaker runtime agents; conversely, using a stronger augments can substantially raise the performance floor for weak agents.

6.7 Takeaways

Our evaluation demonstrates that CodeScout consistently enhances agent performance across diverse models and scaffolds, with structured augmentation outperforming self-augmentation approaches. These findings validate our hypothesis that systematic pre-exploration effectively supplements existing agent architectures. **Effectiveness:** Contextual problem-statement augmentation consistently improves resolution rates and localization across scaffolds and LLMs, with the largest relative gains for weaker runtime agents.

Pipeline vs autonomous augmentation: A separate, structured augmentation pipeline outperforms asking the agent to self-augment during its trajectory.

Design trade-offs: LLM-driven scoping plus relevance filtering is important—simple lexical retrieval (BM25) helps, but LLM scoping yields stronger, more semantically-relevant targets. Augmentation overhead (LLM calls, tokens) is modest relative to downstream gains for most models, though very large-trajectory LLMs maintain higher absolute token budgets.

Deployment guidance: For production settings consider (a) using a moderately capable, cheaper model for augmentation to boost weaker runtime agents, or (b) using a stronger augments to raise the floor of weaker agents when improving the runtime model is not feasible.

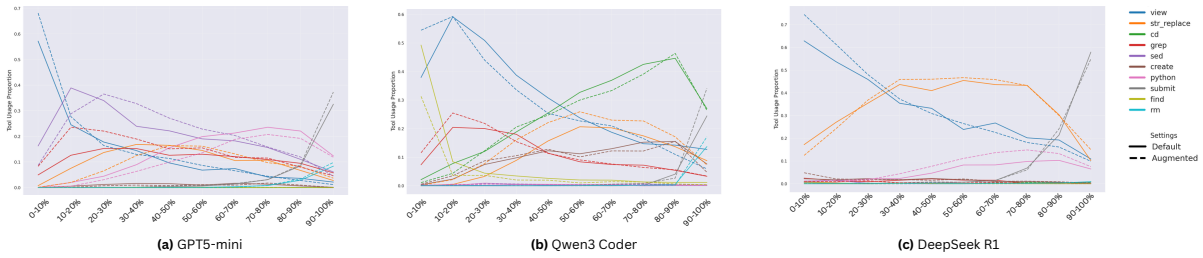


Figure 8: Evolution of tool-call proportions across a trajectory (Default vs CodeScout) for SWE-Agent.

7 Conclusion

We introduced a contextual query refinement framework that transforms vague and underspecified user requests into detailed, actionable problem statements through structured codebase analysis. By combining repository graph construction, high-level context scoping, and multi-perspective code element analysis, our approach systematically grounds problem understanding before attempting solutions. Evaluation on the SWEBench-Verified benchmark demonstrates that this refinement process enhances the effectiveness of AI code assistance, particularly when operating in poorly defined scenarios. Beyond performance gains, our findings highlight the importance of structured problem formulation as a prerequisite for reliable software engineering support. Future work will explore extending this framework to multi-repository settings, integrating deeper program analysis techniques, and applying the approach to collaborative development environments.

8 Limitations

Our evaluation is necessarily constrained by the current state of software engineering benchmarks and tooling. We evaluate CodeScout using three language models (DeepSeek R1, GPT-5-mini, Qwen3-Coder) and three agentic scaffolds (SWE-Agent, OpenHands, Mini-SWE-Agent), which, while representing state-of-the-art systems, constitute only a subset of available approaches. Expanding to additional models and scaffolds would strengthen claims about generalizability but requires substantial computational resources and careful experimental design to ensure fair comparison. Similarly, our analysis focuses exclusively on Python repositories, not by choice but due to the availability of established benchmarks like SWEBench-Verified with reliable evaluation harnesses. Other programming languages lack comparable evaluation infrastructure that provides ground-truth issue

resolution verification at scale, preventing systematic assessment of CodeScout’s effectiveness across language ecosystems with different structural conventions, type systems, and semantic properties. Finally, our evaluation relies on open-source repositories from SWEBench-Verified, which may differ systematically from enterprise codebases in terms of architectural patterns, documentation practices, code quality standards, and proprietary constraints. While open-source repositories enable reproducible research, validating CodeScout’s effectiveness on closed-source or domain-specific industrial code remains important future work that requires industry partnerships and appropriate data sharing agreements.

9 Ethics Statement

Our work focuses on improving AI-assisted software development through enhanced problem statement formulation. We acknowledge the inherent risks of LLM-based code assistance systems, including potential generation of insecure code, introduction of subtle bugs, and perpetuation of biases present in training data. While CodeScout aims to improve agent reliability, users should apply appropriate code review and security practices to any AI-generated solutions. All datasets used in this work—including SWEBench-Verified and associated repositories—are used in accordance with their respective licenses and usage terms. We do not introduce new datasets or collect proprietary code. Our evaluation respects the open-source nature of the benchmark and does not compromise repository security or developer privacy. The computational resources required for our experiments contribute to environmental impact through energy consumption. We encourage practitioners to consider this tradeoff when deploying augmentation systems at scale.

References

- Toufique Ahmed, Kunal Suresh Pai, Premkumar Devanbu, and Earl T Barr. 2024. Automatic semantic augmentation of language model prompts (for code summarization). pages 1–13.
- Shraddha Barke, Michael B. James, and Nadia Polikarpova. 2023. [Grounded copilot: How programmers interact with code-generating models](#). *Proceedings of the ACM on Programming Languages*, 7(OOPSLA1):85–111.
- Islem Bouzenia, Premkumar Devanbu, and Michael Pradel. 2025. Repairagent: An autonomous, llm-based agent for program repair. In *Proceedings of the 47th International Conference on Software Engineering*, ICSE '25, pages 2188–2200. IEEE Computer Society.
- Islem Bouzenia and Michael Pradel. 2025. Understanding software engineering agents: A study of thought-action-result trajectories. *arXiv preprint arXiv:2506.18824*.
- Nan Jiang, Kevin Liu, Thibaud Lutellier, and Lin Tan. 2023. [Impact of code language models on automated program repair](#). In *Proceedings of the 45th International Conference on Software Engineering*, ICSE '23, pages 1430–1442. IEEE Press.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2023. [Swe-bench: Can language models resolve real-world github issues?](#) *Preprint*, arXiv:2310.06770.
- Matthew Jin, Syed Shahriar, Michele Tufano, Xin Shi, Shuai Lu, Neel Sundaresan, and Alexey Svyatkovskiy. 2023. [Inferfix: End-to-end program repair with llms](#). *Preprint*, arXiv:2303.07263.
- Anand Kumar, Vishal Khare, Deepak Sharma, Satyam Kumar, Vijay Saini, Anshul Yadav, Sachendra Jain, Ankit Rana, Pratham Verma, Vaibhav Meena, and Avinash Edubilli. 2025. [Intuition to evidence: Measuring ai's true impact on developer productivity](#). *Preprint*, arXiv:2509.19708.
- Caroline Lemieux, Jeevana Priya Inala, Shuvendu K. Lahiri, and Siddhartha Sen. 2023. [Codamosa: Escaping coverage plateaus in test generation with pre-trained large language models](#). In *Proceedings of the 45th International Conference on Software Engineering*, pages 919–931, Melbourne, Australia. IEEE/ACM.
- 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](#). *Preprint*, arXiv:2409.02977.
- Meili Lu, Xiaobing Sun, Shaowei Wang, David Lo, and Yucong Duan. 2015. Query expansion via wordnet for effective code search. *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, pages 545–549.
- Xiangxin Meng, Zexiong Ma, Pengfei Gao, and Chao Peng. 2024. An empirical study on llm-based agents for automated bug fixing. *arXiv preprint arXiv:2411.10213*.
- Daye Nam, Ahmed Omran, Ambar Murillo, Saksham Thakur, Abner Araujo, Marcel Blistein, Alexander Frömmgen, Vincent Hellendoorn, and Satish Chandra. 2025. Prompting llms for code editing: Struggles and remedies. *arXiv preprint arXiv:2504.20196*.
- Liming Nie, He Jiang, Zhilei Ren, Zeyi Sun, and Xiaochen Li. 2016. Query expansion based on crowd knowledge for code search. *IEEE Transactions on Services Computing*, 9:771–783.
- Stephen Robertson, Hugo Zaragoza, and 1 others. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Francesco Romeo, Luigi Arena, Francesco Blefari, Francesco Aurelio Pironti, Matteo Lupinacci, and Angelo Furfaro. 2025. [Arapaccino: An agentic-rag for policy as code compliance](#). *Preprint*, arXiv:2507.10584.
- Gabriel Ryan, Siddhartha Jain, Mingyue Shang, Shiqi Wang, Xiaofei Ma, Murali Krishna Ramanathan, and Baishakhi Ray. 2024. [Code-aware prompting: A study of coverage-guided test generation in regression setting using llm](#). *Proceedings of the ACM on Software Engineering*, 1(FSE):951–971.
- Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, and 5 others. 2025. [Openhands: An open platform for AI software developers as generalist agents](#). In *The Thirteenth International Conference on Learning Representations*.
- Yanlin Wang, Lianghong Guo, Ensheng Shi, Wenqing Chen, Jiachi Chen, Wanjun Zhong, Menghan Wang, Hui Li, Hongyu Zhang, Ziyu Lyu, and Zibin Zheng. 2023. You augment me: Exploring chatgpt-based data augmentation for semantic code search. *2023 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 14–25.
- Yaqi Wang and Haipei Xu. 2024. [Srsa: A cost-efficient strategy-router search agent for real-world human-machine interactions](#). *Preprint*, arXiv:2411.14574.
- John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan, and Ofir Press. 2024. SWE-agent: Agent-computer interfaces enable automated software engineering. In *Advances in Neural Information Processing Systems*, volume 37.

Zhiqiang Yuan, Mingwei Liu, Shiji Ding, Kaixin Wang, Yixuan Chen, Xin Peng, and Yiling Lou. 2024. *Evaluating and improving ChatGPT for unit test generation*. *Proceedings of the ACM on Software Engineering*, 1(FSE):1703–1726.

Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. 2024. *Autocoderover: Autonomous program improvement*. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*, ISSTA 2024, pages 1592–1604, New York, NY, USA. Association for Computing Machinery.

Albert Ziegler, Eirini Kalliamvakou, Shawn Simister, Ganesh Sittampalam, X. Alice Li, Andrew Rice, Devon Rifkin, and Edward Aftandilian. 2022. *Productivity assessment of neural code completion*. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, MAPS 2022, New York, NY, USA. Association for Computing Machinery.

Output: List of exploration targets $\mathcal{T} = \{(type, name, reasoning)\}$

Stage 1: Exploration Target Identification

You are analyzing a software bug report to identify which files and code components to investigate.

PROBLEM STATEMENT:

{problem_statement}

REPOSITORY STRUCTURE:

{file_tree}

Task: Identify the most promising files, classes, and functions to explore for understanding this issue.

Focus on:

- Files/classes mentioned in the problem statement
- Related components that might be affected
- Common entry points and core functionality
- Test files that might reveal expected behavior

Return exactly 5–10 targets in this format:

1. target_type:file,
target_name:exact_name,
reasoning:why_relevant
2. target_type:class,
target_name:exact_name,
reasoning:why_relevant
3. target_type:function,
target_name:exact_name,
reasoning:why_relevant

Be specific with names—use exact file paths and class/function names from the tree. Each entry must be on a separate line and follow the exact format shown.

Design rationale: We ask the model to produce 5–10 targets to balance coverage with computational cost. The structured output format enables reliable parsing. By requesting explicit reasoning, we encourage the model to articulate its hypotheses about bug locations, which improves target quality.

Appendices

A Prompt Design	11
A.1 Exploration Target Identification . . .	11
A.2 Content Relevance Analysis . . .	11
A.3 Problem Statement Synthesis . . .	12
B Augmentation Algorithm	13
B.1 Main Augmentation Procedure . . .	13
B.2 Computational Complexity	13
C CodeScout: Further Analysis	13
C.1 Qualitative Examples	13
C.2 Quantitative Comparison	13

A Prompt Design

We employ a three-stage prompting strategy to systematically augment problem statements with codebase context. Each prompt is designed to decompose the augmentation task into focused subtasks that elicit specific types of knowledge from the language model. Below we present the complete prompt templates.

A.1 Exploration Target Identification

The first prompt guides the model to identify relevant code locations given a problem statement and repository structure. This stage establishes which parts of the codebase warrant deeper investigation.

Input: Problem statement PS, repository file tree \mathcal{F}

A.2 Content Relevance Analysis

For each exploration target, we retrieve its source code and prompt the model to analyze how it relates to the reported issue. This stage extracts structured insights that will inform the final augmented problem statement.

Input: Problem statement PS, code content c , target metadata θ

Output: Structured analysis with relevance score and insights

Stage 2: Content Relevance Analysis

You are analyzing code content to understand how it relates to a specific bug report. Your analysis will directly feed into an augmented problem statement.

ORIGINAL PROBLEM:

{problem_statement}

CODE CONTENT:

{content}

TARGET INFO: {target.target_type} -
{target.target_name}

Provide analysis in the following format:

RELEVANCE: X/10

ROLE_IN_ISSUE:

Describe how this code relates to the reported issue. What is its role in causing or manifesting the bug?

FIX_LOCATION_HINTS:

Identify specific locations where fixes might be needed (line numbers, methods, conditions). What type of changes might be required (validation, error handling, logic fixes)?

EXPLORATION_INSIGHTS:

Surface key technical details that help understand the issue—algorithms, data flow, dependencies, code patterns, or architectural decisions relevant to the bug. Note test scenarios or edge cases this code suggests should be examined.

ALTERNATIVE_HYPOTHESES:

Consider alternative explanations for the bug based on this code. What related components might also be involved? What different root causes does this analysis suggest?

Note: Focus on insights that will help create actionable exploration hints and fix suggestions. If relevance < 4, keep all sections brief but still provide analysis.

Design rationale: The structured output schema aligns with the final augmented problem statement format. By requesting a numerical relevance score, we enable downstream filtering of low-relevance content. The four semantic categories (role, fix hints, insights, hypotheses) encourage comprehensive analysis while maintaining organization.

A.3 Problem Statement Synthesis

Given the original problem statement and all relevant code analyses, we prompt the model to synthesize a comprehensive augmented problem statement. This stage transforms raw insights into a coherent narrative that guides developers.

Input: Original problem statement PS, filtered analyses \mathcal{A}_{rel}

Output: Augmented problem statement PS'

Stage 3: Problem Statement Synthesis

You are creating an enhanced problem statement that will help developers quickly understand and fix a software issue.

ORIGINAL PROBLEM STATEMENT:
{original_ps}

CODE ANALYSIS RESULTS: For each analyzed target:

{i}. **TARGET:** {target_type} - {target_name}
RELEVANCE: {score}/10
ROLE IN ISSUE: {role_description}
FIX LOCATION HINTS: {hints}
EXPLORATION INSIGHTS: {insights}
ALTERNATIVE HYPOTHESES: {hypotheses}

Instructions: Synthesize these analyses into a comprehensive problem statement. Include all relevant details from the original statement. Use the code analysis as context to inform your narrative—don't simply copy it verbatim.

Structure your output with these sections:

Issue Description: Rewrite the core problem clearly, incorporating technical insights.

Reproduction Steps: Include original steps if available. If the original issue includes code, preserve it here—this is essential. Add technical details about what happens internally. Include relevant stack traces or error patterns. Suggest additional reproduction scenarios if helpful.

Expected Behavior: Describe what should happen instead, informed by code analysis.

Exploration Hints:

- *Files to Examine:* file1.py - why relevant (role in issue, key findings)

- *Key Classes/Functions:* ClassName.method() - what to look for

- *Areas of Interest:* Specific code patterns or logic to investigate

Fix Hints:

- *High Confidence Locations:* Where fixes are likely needed and why

- *Implementation Suggestions:* Concrete fix ideas with limitations

- *Alternative Hypotheses:* Other potential root causes

Keep technical language precise but accessible. Focus on actionable insights.

Design rationale: The detailed template structure ensures consistency across augmented problem statements. By explicitly requesting preservation of reproduction code and technical details, we maintain the diagnostic value of the original statement while enriching it with codebase context. The hierarchical organization mirrors how developers naturally think about debugging: understanding the issue, knowing where to look, and hypothesizing about fixes.

Algorithm 1 Problem Statement Augmentation

```
1: Input: Problem statement PS, codebase  $\mathcal{C}$ , language model  $\mathcal{M}$ , threshold  $\tau$ 
2: Output: Augmented problem statement PS'
3:
4: // Stage 1: Exploration Target Identification

5:  $\mathcal{F} \leftarrow \text{GETFILETREE}(\mathcal{C})$ 
6:  $p_1 \leftarrow \text{FORMATPROMPT}_{\text{explore}}(\text{PS}, \mathcal{F})$ 
7:  $r_1 \leftarrow \mathcal{M}(p_1)$ 
8:  $\mathcal{T} \leftarrow \text{PARSETARGETS}(r_1)$ 
9:
10: // Stage 2: Content Retrieval and Analysis
11:  $\mathcal{A} \leftarrow []$ 
12: for each target  $\theta \in \mathcal{T}$  do
13:   content  $\leftarrow \text{RETRIEVECODE}(\theta, \mathcal{C})$ 
14:   if content successfully retrieved then
15:      $p_2 \leftarrow \text{FORMATPROMPT}_{\text{analyze}}(\text{PS}, \text{content}, \theta)$ 
16:      $r_2 \leftarrow \mathcal{M}(p_2)$ 
17:      $a \leftarrow \text{PARSEANALYSIS}(r_2)$ 
18:      $\mathcal{A} \leftarrow \mathcal{A} \cup \{a\}$ 
19:   end if
20: end for
21:
22: // Stage 3: Relevance Filtering
23:  $\mathcal{A}_{\text{rel}} \leftarrow \{a \in \mathcal{A} \mid a.\text{relevance} \geq \tau\}$ 
24:
25: // Stage 4: Synthesis
26: insights  $\leftarrow \text{FORMATINSIGHTS}(\mathcal{A}_{\text{rel}})$ 
27:  $p_3 \leftarrow \text{FORMATPROMPT}_{\text{synthesize}}(\text{PS}, \text{insights})$ 

28: PS'  $\leftarrow \mathcal{M}(p_3)$ 
29:
30: return PS'
```

B Augmentation Algorithm

Our approach processes each problem statement through a structured pipeline that leverages language models to systematically gather and synthesize codebase context. Algorithm 1 presents the core augmentation procedure.

B.1 Main Augmentation Procedure

Key design choices: We employ a **relevance threshold** τ to filter analyses before synthesis, ensuring only high-quality insights inform the final augmented statement. The pipeline gracefully handles retrieval failures by skipping unavailable targets rather than halting, making it robust to incom-

plete code references.

B.2 Computational Complexity

LLM API Calls: Each problem statement requires $2 + |\mathcal{T}_{\text{valid}}|$ language model calls:

- 1 call for exploration target identification
- $|\mathcal{T}_{\text{valid}}|$ calls for content analysis (typically 5–10 targets)
- 1 call for synthesis

For a dataset of size n , the total number of LLM calls is $O(n \cdot |\mathcal{T}|)$. In practice, with approximately 7 valid targets per instance, we observe roughly 9 API calls per problem statement.

C CodeScout: Further Analysis

C.1 Qualitative Examples

We present detailed visualizations of the CodeScout context retrieval pipeline across five diverse instances from the SWE-bench Verified dataset. Each instance demonstrates the three-stage pipeline: (1) high-level context scoping to identify relevant files, classes, and functions, (2) fine-grained content analysis with relevance scoring, and (3) problem statement augmentation. We show results from three different retrieval methods using DeepSeek-R1, Qwen3-Coder, and GPT-5-mini.

C.2 Quantitative Comparison

We present comprehensive comparative analysis of the CodeScout pipeline for the chosen LLMs: DeepSeek-R1, Qwen3-Coder, and GPT-5-mini.

Score Distributions. Figure 24 shows the distribution of relevance scores assigned by each method across all targets. Qwen3-Coder and GPT-5-mini show higher concentrations of scores in the 7-9 range, while DeepSeek-R1 exhibits a broader distribution with more scores in the lower ranges. Figure 25 displays the distribution of maximum scores per instance, showing that Qwen3-Coder and GPT-5-mini more frequently assign maximum scores of 8-9, whereas DeepSeek-R1 shows greater variability.

Target Coverage. Figure 26 presents the distribution of the number of targets identified per instance for each method. DeepSeek-R1 tends to identify more targets per instance with a broader distribution, while Qwen3-Coder shows a tighter distribution around 6 targets. Figure 27 illustrates

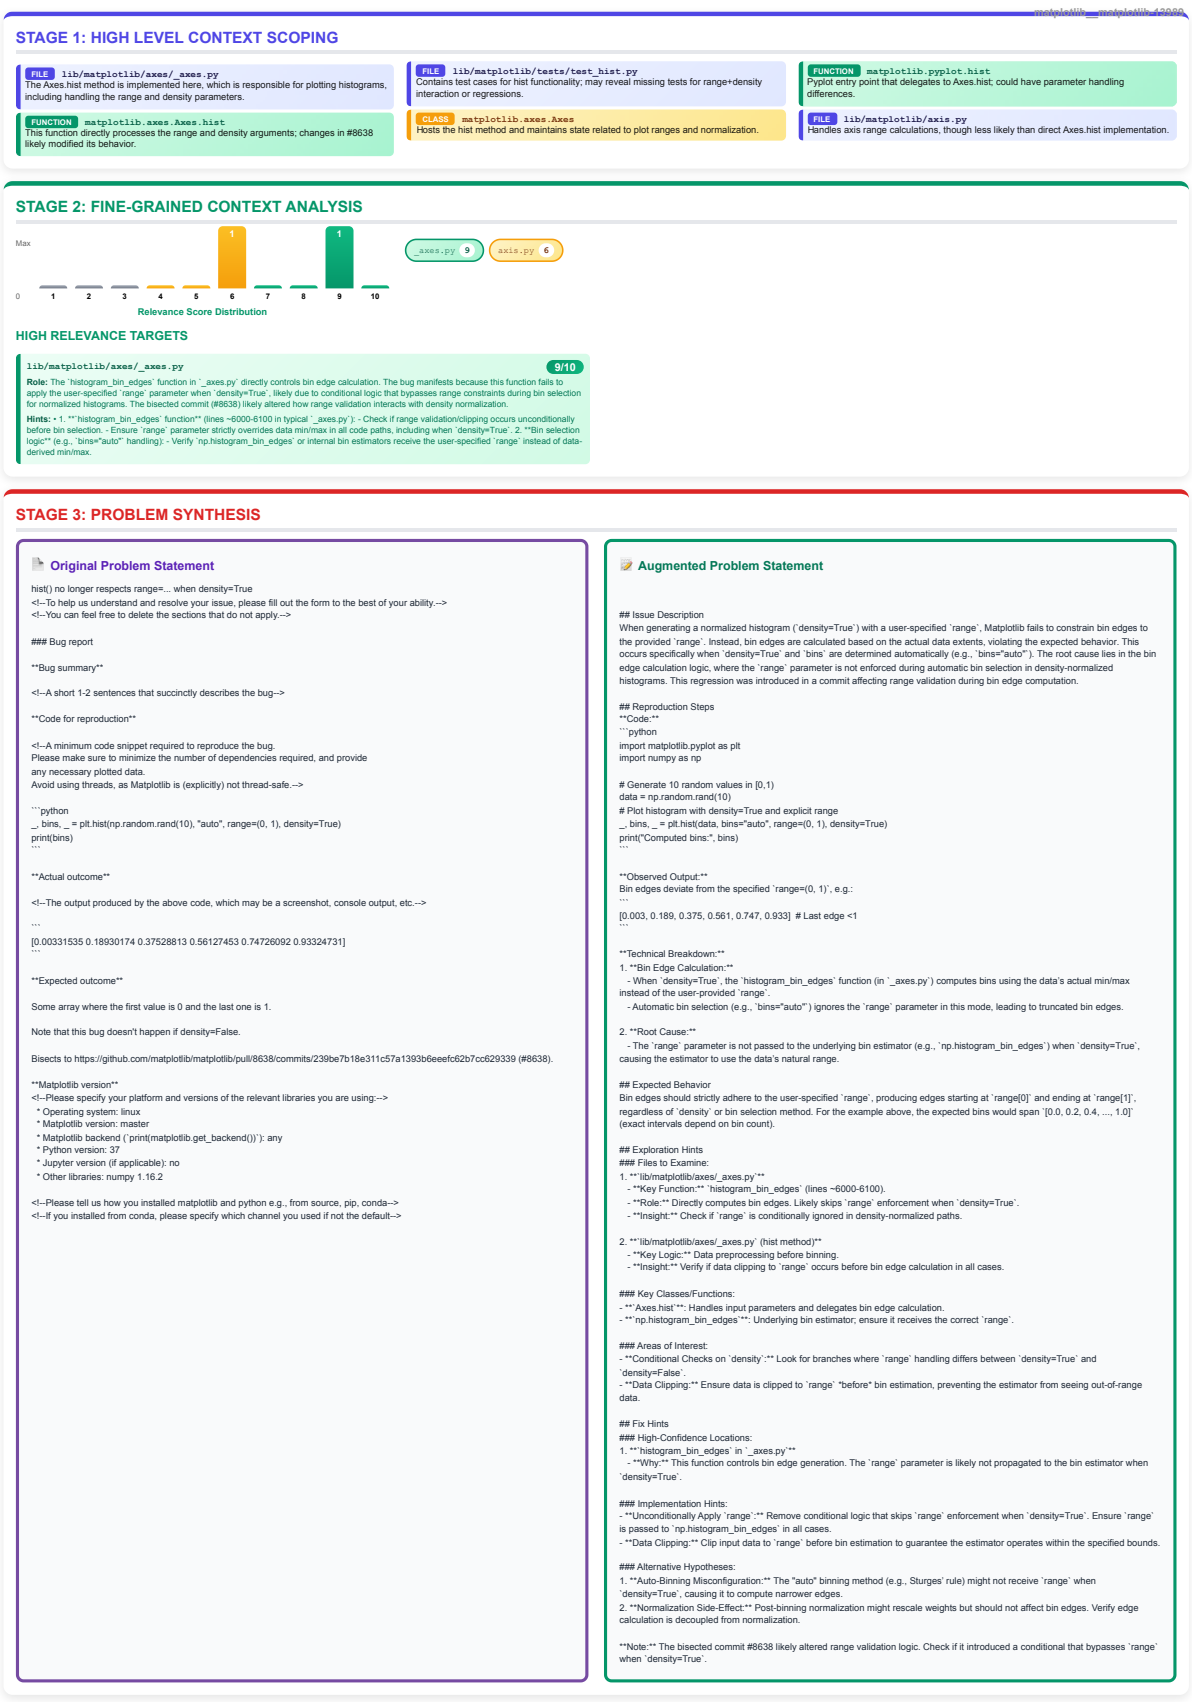


Figure 9: CodeScout visualization for matplotlib__matplotlib-13989 with DeepSeek-R1.

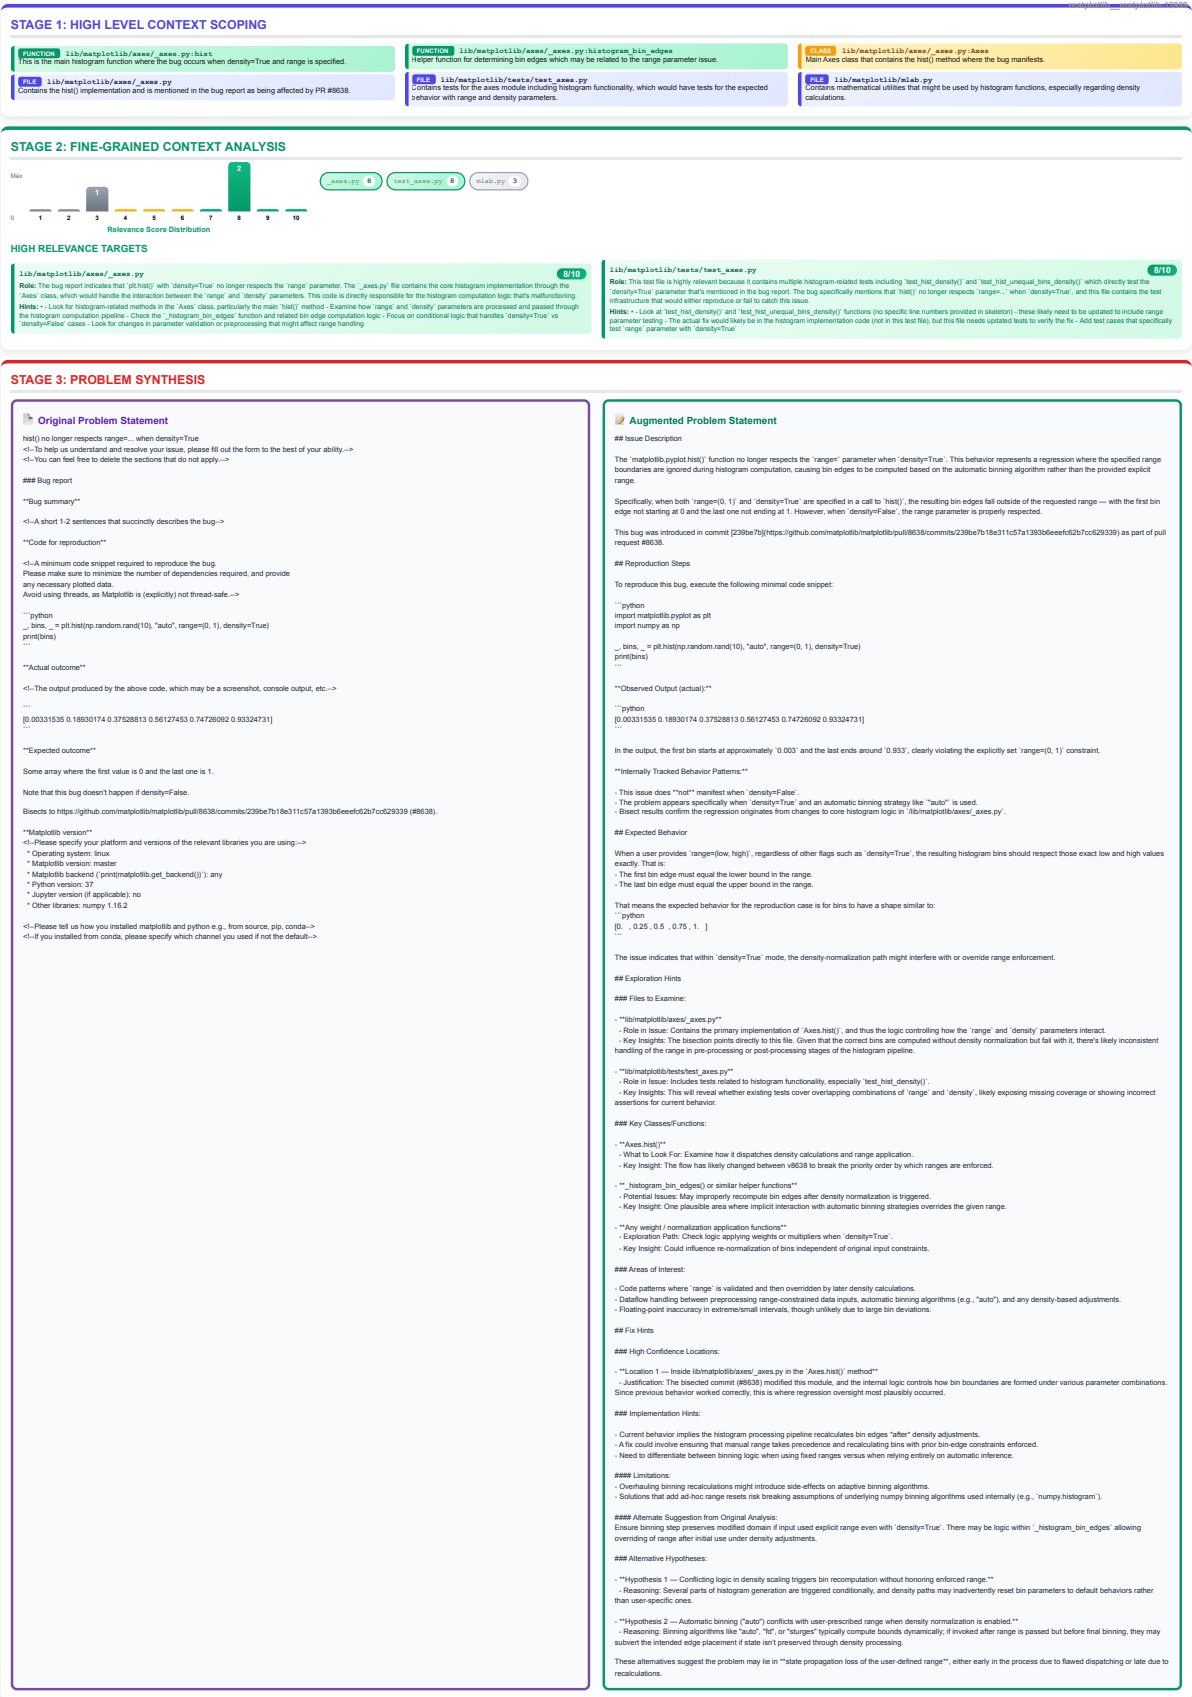


Figure 10: CodeScout visualization for matplotlib__matplotlib-13989 with Qwen3-Coder.

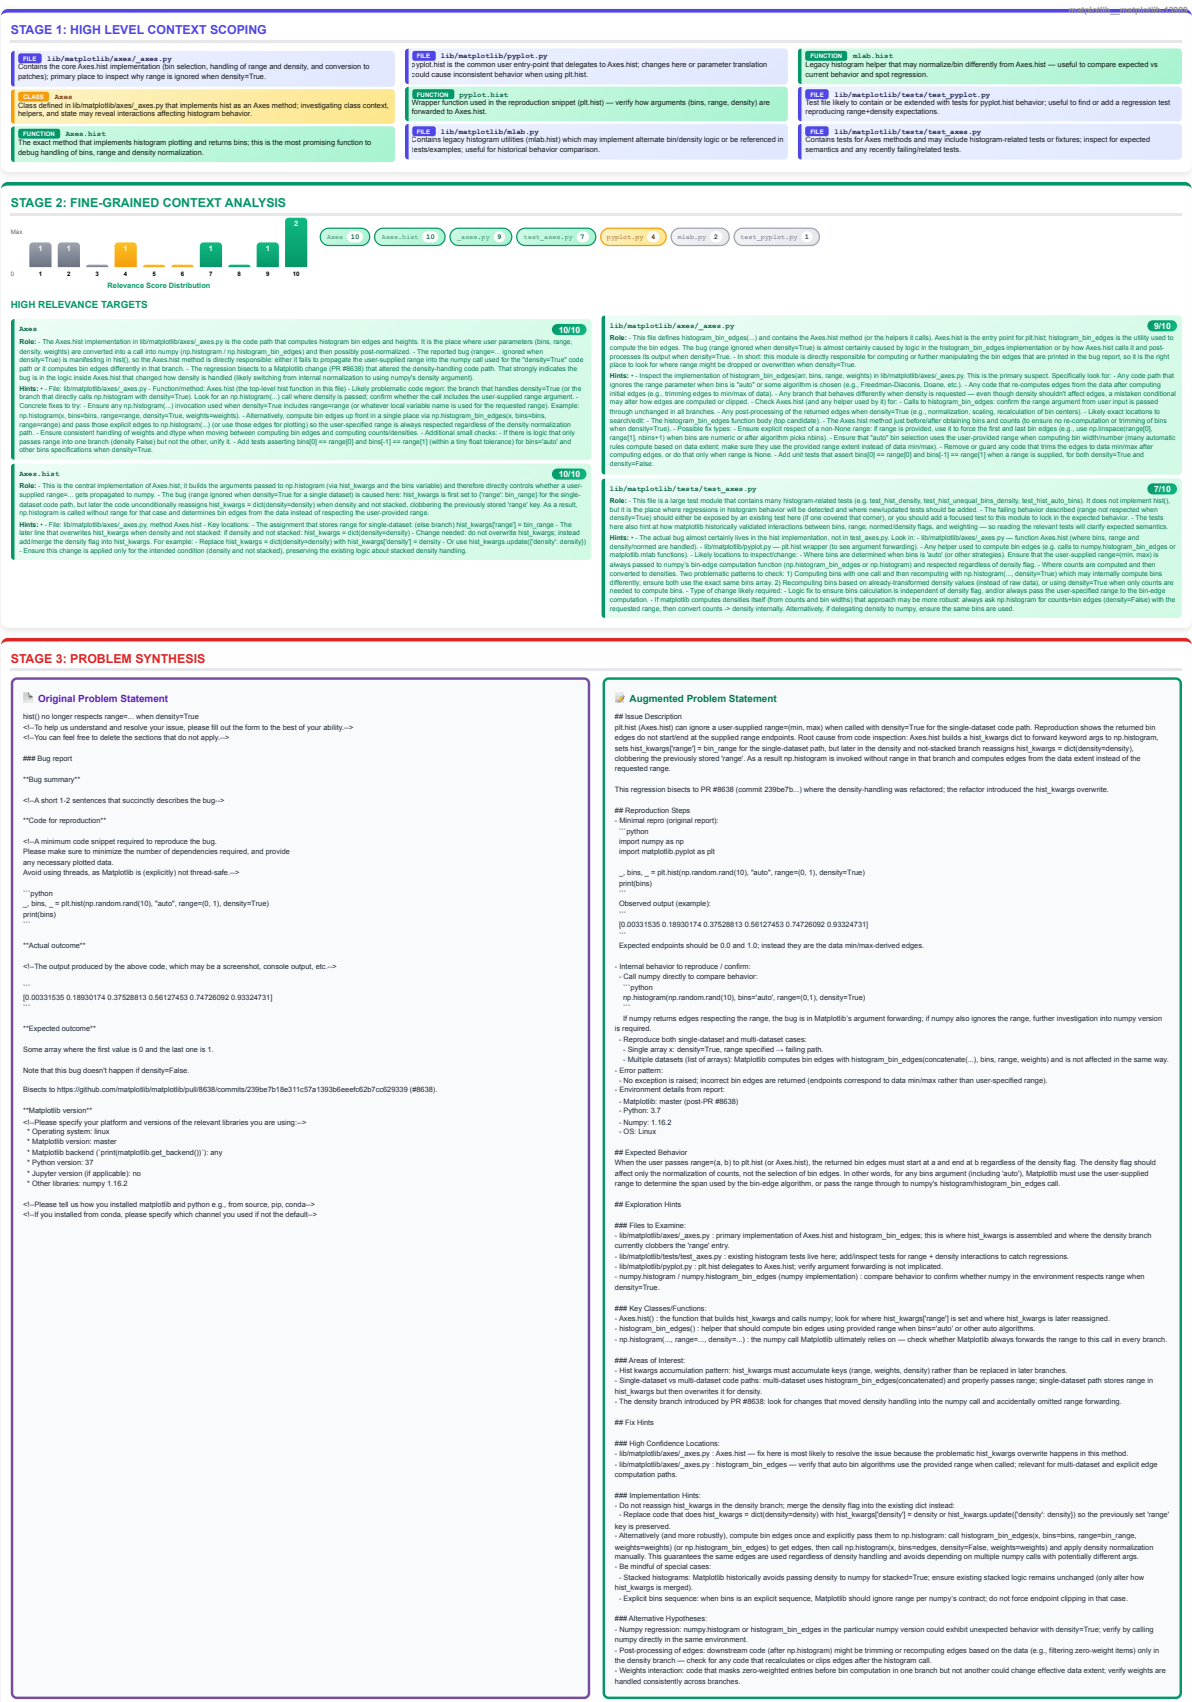


Figure 11: CodeScout visualization for matplotlib__matplotlib-13989 with GPT-5-mini.

STAGE 1: HIGH LEVEL CONTEXT SCOPING

FILE `src/_pytest/logging.py`
Contains `LogCaptureFixture` (pytest) implementation, including `clear()` and `get_records()` methods mentioned in the problem.

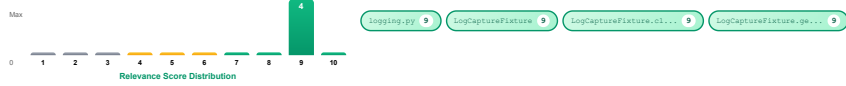
CLASS `LogCaptureFixture`
The `caplog` fixture instance; manages record storage and clearing, directly involved in the conflict.

FUNCTION `LogCaptureFixture.get_records`
Returns records filtered by test phase but references outdated list after `clear()`.

FUNCTION `LogCaptureFixture.clear`
Replaces records list instead of clearing it, causing divergence from `get_records()`.

FUNCTION `LogCaptureFixture.setup`
Initializes `get_records` to same list as `records` during test setup (line 699 in problem links).

STAGE 2: FINE-GRAINED CONTEXT ANALYSIS



HIGH RELEVANCE TARGETS

FILE `src/_pytest/logging.py` **9/10**
Role: The `LogCaptureFixture`'s `clear()` method replaces the `records` list with a new empty list instead of clearing the existing one. This causes `get_records()` (which references the original list) to become decoupled from `caplog.records` after `clear()`. The divergence occurs because `get_records()` and `records` are initialized to share the same list during test setup, but `clear()` breaks this relationship by reassignment.
Hints: 1. `LogCaptureFixture.clear()` method: Replace `self.records = []` with `self.records.clear()` to mutate the existing list instead of creating a new one. 2. `LogCaptureFixture.get_records()` implementation: Verify it dynamically accesses the current `records` list rather than caching an old reference.

FUNCTION `LogCaptureFixture.clear` **9/10**
Role: The `LogCaptureFixture.clear()` method calls `self.handler.reset()`, which replaces the handler's internal records list rather than clearing it in place. This causes `caplog.records` (which references `handler.records`) to point to a new list, while `get_records()` retains a reference to the original pre-clear list from test setup. This divergence creates the frozen state observed in the bug.
Hints: 1. Modify `LogCaptureFixture.clear()` to update `self.records` after resetting the handler (e.g., `self.records = self.handler.records.post_reset`). 2. Change the handler's `reset` logic to mutate the existing list (e.g., `self.records.clear()`) instead of replacing it. 3. Ensure `get_records()` dynamically references the current handler records rather than caching an initial reference.

FUNCTION `LogCaptureFixture.get_records` **9/10**
Role: The `LogCaptureFixture`'s `clear()` method replaces the `records` list with a new empty list instead of clearing the existing list. This causes `get_records()` (which references the original list) to diverge from `caplog.records` after `clear()`, as they now point to different lists. The core conflict arises from list identity vs mutation semantics.
Hints: 1. `LogCaptureFixture.get_records()` method: Replace `self.records = []` with `self.records.clear()` to mutate the existing list instead of creating a new one (preserves reference consistency). 2. `LogCaptureFixture.clear()` method: Ensure `clear()` updates all relevant phase lists or uses in-place clearing. 3. `LogCaptureFixture.get_records()` implementation: Verify `LogCaptureHandler`'s internal buffer is also cleared in-place to maintain consistency with `caplog.records`.

FUNCTION `LogCaptureFixture.get_records` **9/10**
Role: The `get_records()` method directly accesses records stored in the test item's stash (`self._item_stash`), which is initially synchronized with `caplog.records`. However, when `clear()` replaces `caplog.records` with a new list instead of clearing it in place, the stash retains a reference to the old list. This causes `get_records()` to return stale data after `clear()`, directly manifesting the divergence described in the bug.
Hints: 1. `LogCaptureFixture.get_records()` method: Replace `self.records = []` with in-place clearing (e.g., `self.records.clear()`) to preserve the original list reference. 2. `LogCaptureFixture.get_records()` implementation: Ensure the stash and `caplog.records` always reference the same list object, even after `clear()`.

STAGE 3: PROBLEM SYNTHESIS

```
Original Problem Statement
caplog.get_records and caplog.clear conflict
# Description
caplog.get_records() gets decoupled from actual caplog records when caplog.clear() is called. As a result, after caplog.clear() is called, caplog.get_records() is frozen: it does not get cleared, nor does it get new records.

During test set up it is [set to the same list]https://github.com/pytest-dev/pytest/blob/28e8c852e947704655a3c32d5718483133869fsrc/_pytest/logging.py#L699 as 'caplog.records', but the latter gets [replaced rather than cleared]https://github.com/pytest-dev/pytest/blob/28e8c852e947704655a3c32d5718483133869fsrc/_pytest/logging.py#L345 in 'caplog.clear()', which diverges the two objects.

# Reproducible example
'''python
import logging

def test(caplog) -> None:
    def verify_consistency() -> None:
        assert caplog.get_records("call") == caplog.records

    verify_consistency()
    logging.warning("test")
    verify_consistency()
    caplog.clear()
    verify_consistency() # fails: assert [<LogRecord ...>, 8, 'test'] == []

# Environment details
Arch: Linux, Python 3.9.10:
...
Package Version
-----
attrs 21.4.0
iniconfig 1.1.1
packaging 21.3
pip 22.0.4
pluggy 1.0.0
py 1.11.0
pyrsistent 3.0.0
pytest 7.1.1
setuptools 60.10.0
tomli 2.0.1
wheel 0.37.1
...

Augmented Problem Statement
## Issue Description
When caplog.clear() is called, it replaces the internal 'records' list with a new empty list instead of clearing the existing one. This breaks referential consistency between 'caplog.records' and 'caplog.get_records()', which initially share the same list. After 'clear()', 'get_records()' continues to reference the original list, leading to stale data and frozen behavior. The root cause lies in mutation-vs-reassignment semantics in 'LogCaptureFixture' state management, exacerbated by phase-specific record tracking in pytest's logging infrastructure.

## Reproduction Steps
'''Code'''
'''python
import logging

def test(caplog) -> None:
    def verify_consistency() -> None:
        assert caplog.get_records("call") == caplog.records

    verify_consistency() # Passes: both are empty
    logging.warning("test")
    verify_consistency() # Passes: both contain the new log
    caplog.clear()
    verify_consistency() # Fails: get_records() retains pre-clear logs

## Internal Behavior:
1. During test setup, 'caplog.records' and 'get_records("call")' are initialized to reference the same list.
2. 'caplog.clear()' replaces 'self.records' with a new empty list, while 'get_records()' retains the original reference.
3. Subsequent logging appends to the new 'records' list, but 'get_records()' continues to read from the old, now-stale list.

## Error Pattern:
- Post-'clear()' assertions fail with messages like 'assert [<LogRecord ...>] == []', indicating divergent record states.

## Expected Behavior
After calling 'caplog.clear()':
1. Both 'caplog.records' and 'caplog.get_records()' should return an empty list.
2. New logs added after 'clear()' should appear in both 'records' and 'get_records()'.
3. The internal list reference shared between these properties should remain consistent across all test phases.

## Exploration Hints
### Files to Examine:
1. src/_pytest/logging.py
- Role: Contains LogCaptureFixture and its clear()/get_records() methods.
- Key Insight: clear() replaces self.records with a new list, breaking synchronization with get_records().

### Key Classes/Functions:
1. LogCaptureFixture.clear()
- Issue: Uses self.records = [] instead of in-place self.records.clear().
- Impact: Reassignment decouples records from get_records(), which retains the old list.
2. LogCaptureHandler.reset()
- Suspicion: May replace the handler's internal buffer list, propagating inconsistency to caplog.records.

### Areas of Interest:
- List Identity vs. Mutation: Verify whether all record storage uses the same list instance.
- Phase-Specific Tracking: Check if setup/call/teardown phases cache separate list references.

## Fix Hints
### High Confidence Locations:
1. LogCaptureFixture.clear()
- Why: Directly responsible for replacing self.records instead of mutating it.
- Fix: Replace self.records = [] with self.records.clear().

### Implementation Hints:
1. In-Place List Clearing
- Modify clear() to mutate the existing records list:
'''python
def clear(self) -> None:
    self.records.clear() # Instead of self.records = []
    self.handler.reset() # Ensure handler also clears in-place
...
- Limitation: Requires LogCaptureHandler.reset() to also clear its buffer without reassignment.
2. Handler Synchronization
- Update LogCaptureHandler.reset() to use self.records.clear() instead of self.records = [].

### Alternative Hypotheses:
1. Phase-Specific Caching
- Possibility: get_records(when) caches phase-specific lists not reset by clear().
- Investigate: Whether phase lists (e.g., setup/call/teardown) share the same reference as caplog.records.
2. Stash Reference Staleness
- Possibility: get_records() pulls from self._item_stash, which isn't updated after clear().
- Check: If the stash synchronizes with self.records dynamically or caches an initial reference.
```

Figure 12: CodeScout visualization for `pytest-dev__pytest-10051` with DeepSeek-R1.

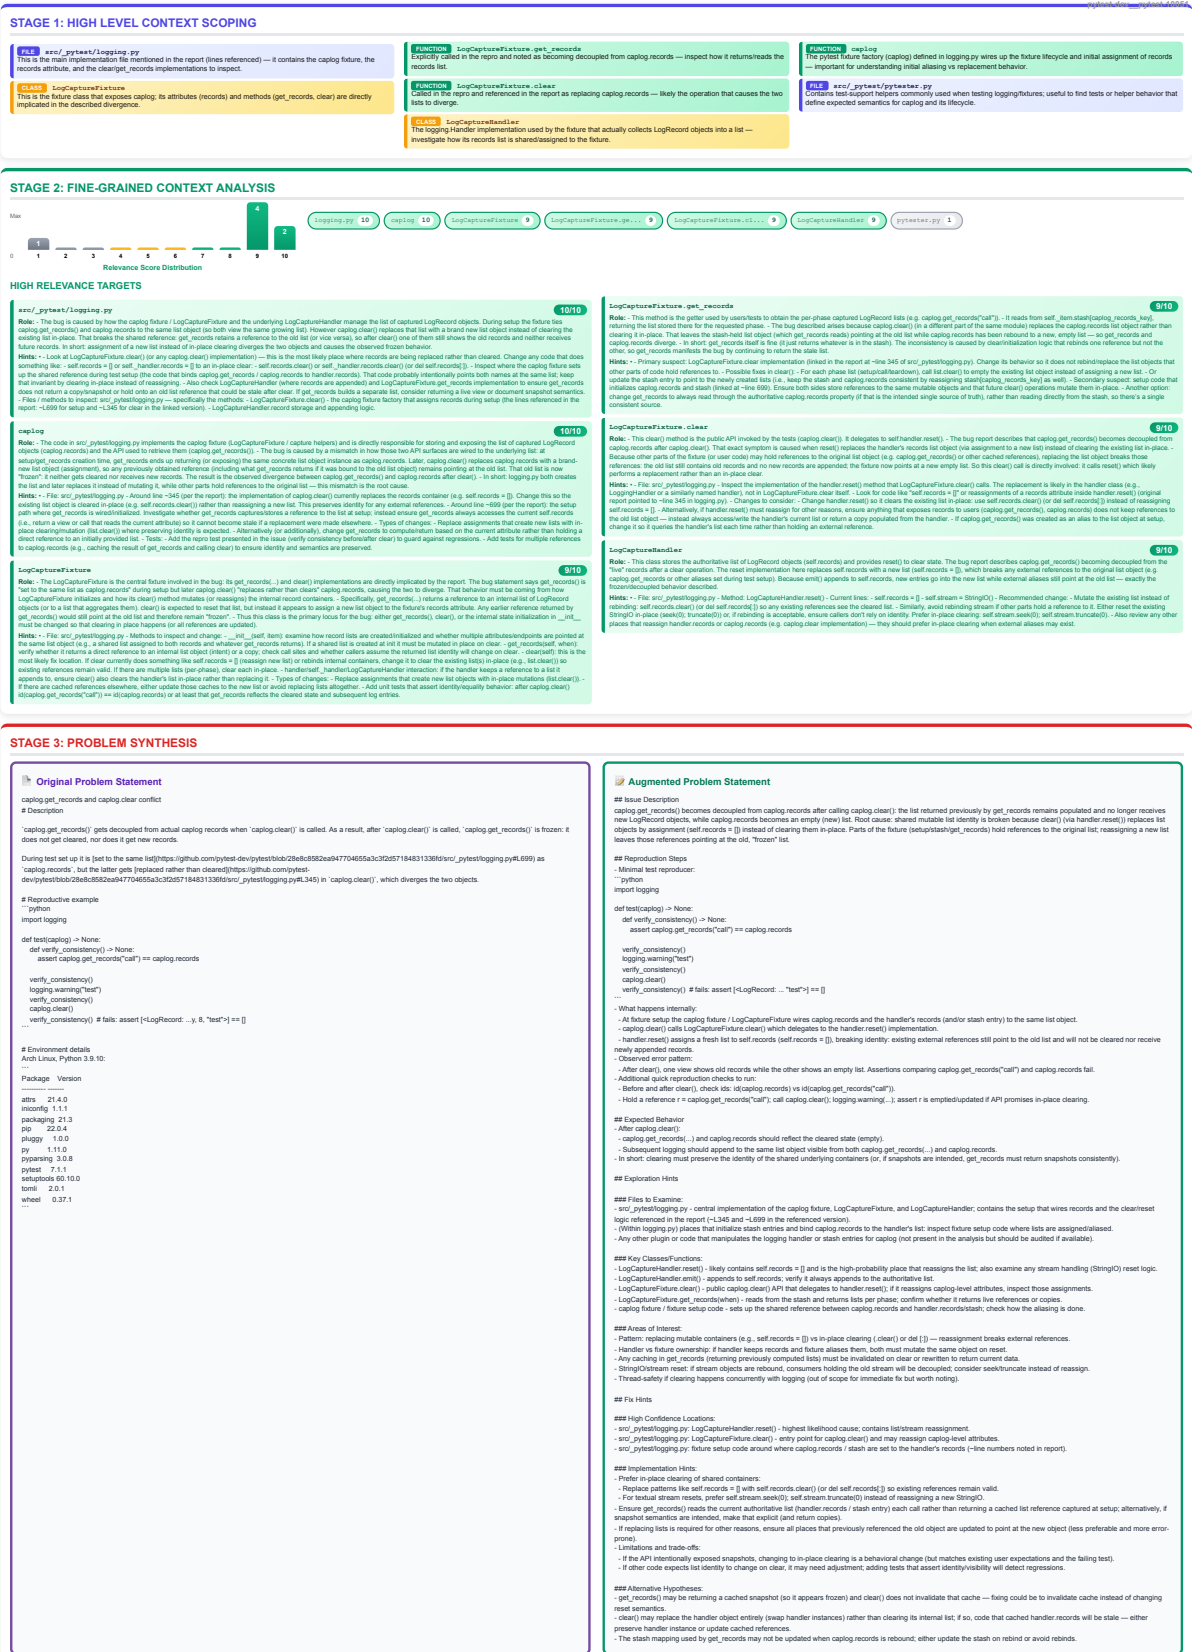


Figure 14: CodeScout visualization for pytest-dev__pytest-10051 with GPT-5-mini.

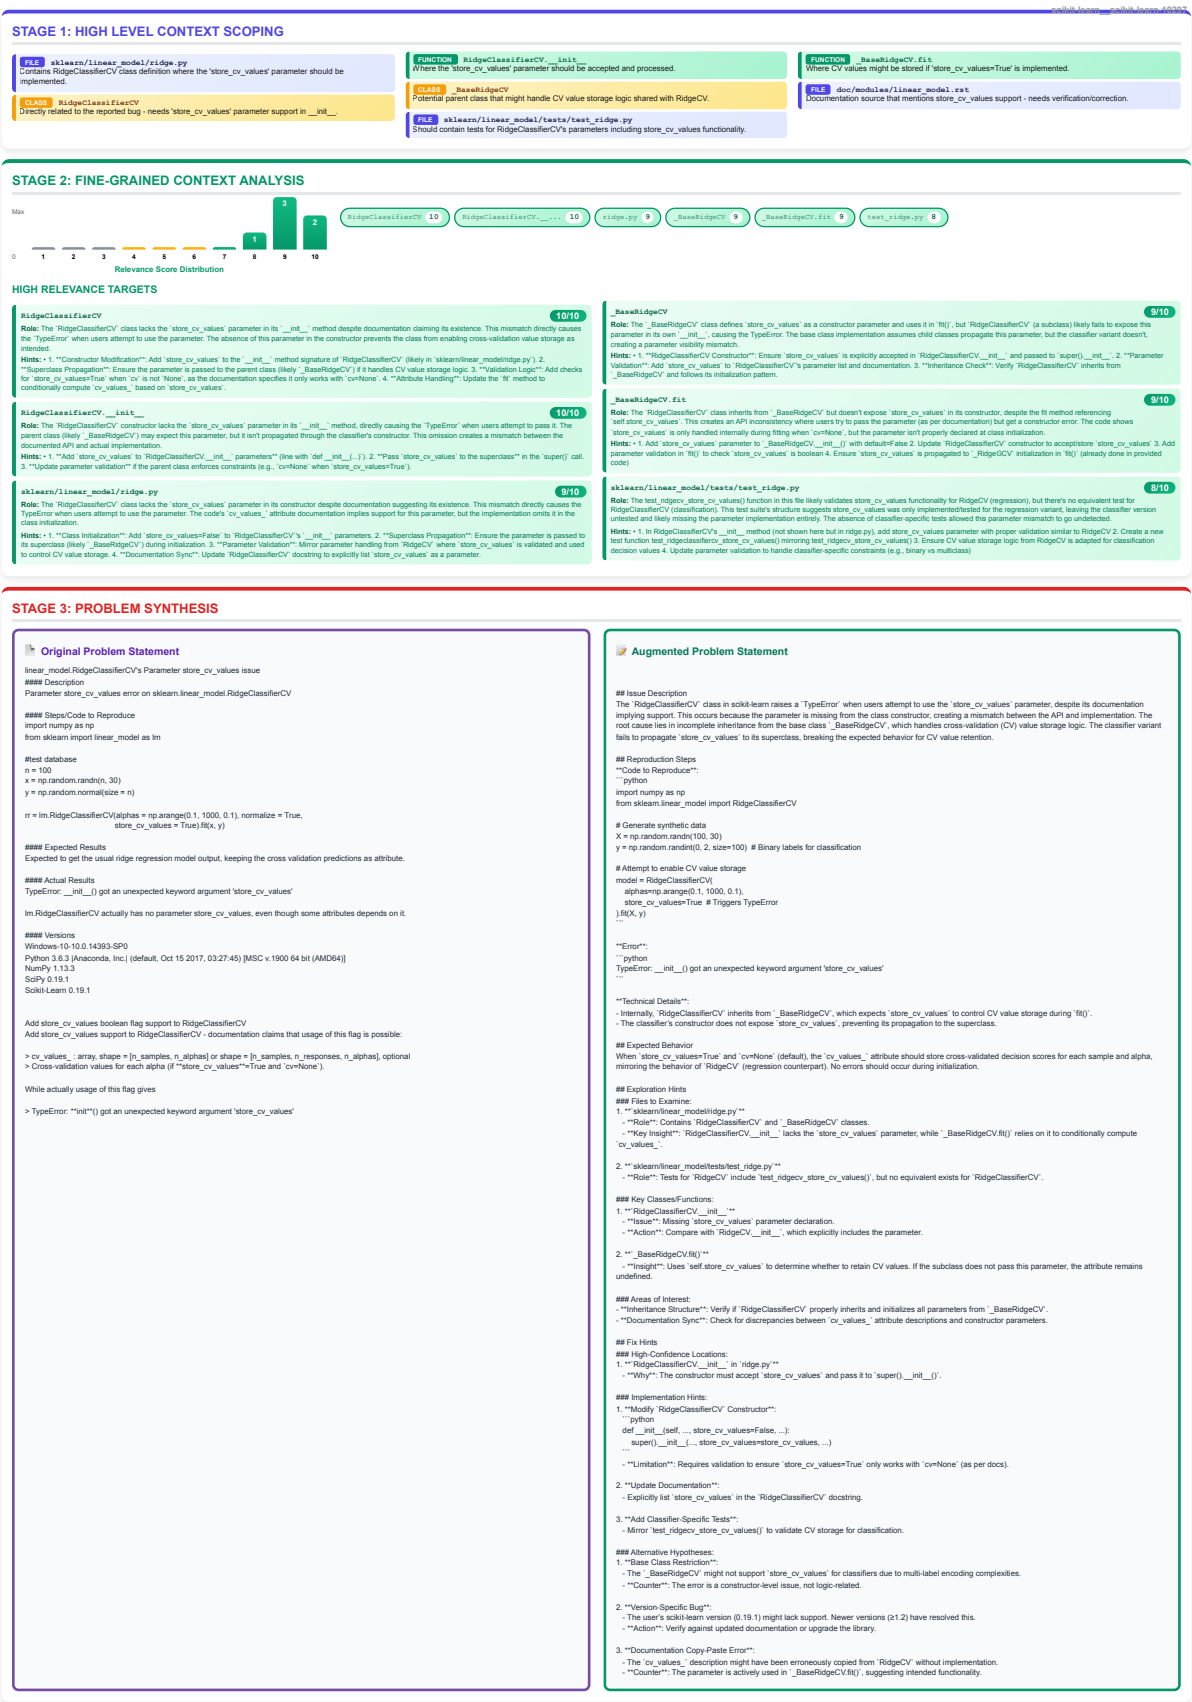


Figure 15: CodeScout visualization for scikit-learn__scikit-learn-10297 with DeepSeek-R1.

STAGE 1: HIGH LEVEL CONTEXT SCOPING

Issue: sklearn/linear_model/ridge.py: This file contains the RidgeClassifierCV class implementation where the store_cv_values parameter is missing from the constructor.

Class: RidgeClassifierCV: The main class where the store_cv_values parameter needs to be added to match the documentation and user expectations.

File: sklearn/linear_model/feature_selection.py: This file contains tests for RidgeClassifierCV and would need to be updated to test the new store_cv_values functionality.

Class: RidgeCV: RidgeClassifierCV likely inherits or shares functionality with RidgeCV, which may already have store_cv_values implemented.

Function: RidgeClassifierCV.__init__: The constructor method that needs to be modified to accept the store_cv_values parameter.

Function: RidgeClassifierCV.fit: The fit method where the store_cv_values logic would be implemented to store cross-validation values when requested.

STAGE 2: FINE-GRAINED CONTEXT ANALYSIS

Relevance Score Distribution: A bar chart showing relevance scores for various files. The highest scores are for sklearn/linear_model/ridge.py (score 9) and sklearn/linear_model/feature_selection.py (score 8).

High Relevance Targets:

- sklearn/linear_model/ridge.py: The code in sklearn/ridge.py directly refers to the reported issue because RidgeClassifierCV class is defined in this file and is missing the store_cv_values parameter in its constructor, despite having documentation that references the parameter and the cv_values attribute. The issue manifests because the parameter is not accepted during initialization, even though the underlying functionality appears to be partially implemented as evidenced by the documented cv_values attribute. Note: 1. "Constructor signature": Add store_cv_values parameter to __init__ method 2. "RidgeClassifierCV class signature": Update the parameter list to include store_cv_values=False 3. "Parent class BaseRidgeCV signature": Ensure the parameter is properly passed to parent classes that handle the actual CV value storage logic 4. "Documentation consistency": Verify that the parameter is properly documented in the class docstring.
- RidgeClassifierCV: The code directly refers to the reported issue because the RidgeClassifierCV class __init__ method does not accept a store_cv_values parameter, even though the class documentation mentions the cv_values attribute that should be populated when this flag is set to True. The mismatch between documented API and actual implementation causes the TypeError when users try to use the parameter. Note: 1. "Constructor signature": Add store_cv_values parameter to __init__ method 2. "Constructor implementation": Store the parameter as instance variable 3. "Fit method": Modify CV logic to handle storing CV values when the flag is True and cv=None 4. "Documentation consistency": Ensure docstring accurately reflects implemented functionality.
- RidgeClassifierCV.__init__: The code in the root case of the bug. The __init__ method of RidgeClassifierCV explicitly does not include store_cv_values as a parameter, even though the documentation claims it should be supported and the parent class __init__ method does not raise an error. This mismatch between documented API and actual implementation causes the TypeError when users try to use the parameter. Note: 1. Line 1: Add store_cv_values=False parameter to the __init__ method signature. Line 3: Pass store_cv_values=True to the parent class __init__. Consider adding validation logic to ensure store_cv_values only works when cv=None is documented.
- sklearn/linear_model/feature_selection.py: The file has a already relevant to the reported issue as it contains tests for RidgeCV functionality, including a specific test for store_cv_values support (test_ridgecv_store_cv_values). The absence of equivalent tests for RidgeClassifierCV's store_cv_values parameter suggests incomplete implementation coverage. The code results that while RidgeCV supports store_cv_values, RidgeClassifierCV does not, which directly manifests the bug reported when users expect the parameter to be available. Note: 1. Look at test_ridgecv_store_cv_values function to understand expected behavior and implementation pattern. Check RidgeClassifierCV class definition (likely in ridge.py) for missing store_cv_values parameter in __init__. 2. Examine how cv_values attribute is computed and stored in RidgeCV to replicate in RidgeClassifierCV. Add equivalent test function for RidgeClassifierCV's store_cv_values functionality.
- RidgeCV: The code defines the RidgeCV class which properly implements the store_cv_values parameter. The issue reported is about RidgeClassifierCV missing the parameter, while this class shows the correct implementation in the parent class. The code is relevant because it demonstrates the expected behavior and parameter support that should also exist in RidgeClassifierCV. The bug manifests because RidgeClassifierCV doesn't expose this parameter while its documentation suggests it should be available. Note: 1. Need to examine RidgeClassifierCV class definition (likely in the same file or nearby). Look for BaseRidgeCV class to understand shared functionality. Check __init__ method of RidgeClassifierCV for missing parameter. Examine how cv_values attribute is handled in both classes. Look at the class inheritance structure and parameter passing.
- RidgeClassifierCV.fit: The code shows the fit method of RidgeClassifierCV, which delegates to BaseRidgeCV. The issue is that RidgeClassifierCV doesn't expose the store_cv_values parameter in its constructor, even though its parent class BaseRidgeCV expects it and the documentation mentions it. The code is directly involved in the bug because when the parameter should be passed through but isn't being passed due to the missing constructor parameter. Note: 1. Check the RidgeClassifierCV.__init__ method (not shown in this snippet) to add store_cv_values parameter. Line where BaseRidgeCV.__init__(self, X, y, sample_weight=weights, cv=None) may need to pass through the store_cv_values parameter. The BaseRidgeCV class implementation (not shown) where the actual store_cv_values logic is handled.
- RidgeClassifierCV.__init__: The code shows the __init__ method of RidgeClassifierCV, which delegates to BaseRidgeCV. The issue is that RidgeClassifierCV doesn't expose the store_cv_values parameter in its constructor, even though its parent class BaseRidgeCV expects it and the documentation mentions it. The code is directly involved in the bug because when the parameter should be passed through but isn't being passed due to the missing constructor parameter. Note: 1. Check the RidgeClassifierCV.__init__ method (not shown in this snippet) to add store_cv_values parameter. Line where BaseRidgeCV.__init__(self, X, y, sample_weight=weights, cv=None) may need to pass through the store_cv_values parameter. The BaseRidgeCV class implementation (not shown) where the actual store_cv_values logic is handled.

STAGE 3: PROBLEM SYNTHESIS

Original Problem Statement: sklearn/linear_model/ridge.py Parameter store_cv_values issue. Description: Parameter store_cv_values error on sklearn.linear_model.RidgeClassifierCV. Steps to Reproduce: Import numpy as np, from sklearn import linear_model as lm, #use database, n = 100, x = np.random.randn(30), y = np.random.normal(size=n), fit = lm.RidgeClassifierCV(alpha = np.arange(0.1, 1000.0, 1), normalize = True, store_cv_values=True)(x, y). Expected Results: Expected to get the usual ridge regression model output, keeping the cross validation predictions as attribute. Actual Results: TypeError: __init__() got an unexpected keyword argument 'store_cv_values'. In RidgeClassifierCV actually has no parameter store_cv_values, even though some attributes depends on it. Versions: Windows 10 20.0.13430.520, Python 3.8.3 (tags/branchpoint: 3.8.3), numpy 1.19.3, SciPy 1.5.1, Scikit-learn 0.19.1. Add store_cv_values boolean flag support to RidgeClassifierCV. Add store_cv_values support to RidgeClassifierCV - documentation claims that usage of this flag is possible: cv_values: array, shape = (n_samples, n_alpha) or shape = (n_samples, n_responses, n_alpha), optional - Cross-validation values for each alpha (if 'ridge_cv_values=True and cv=None'). While actually usage of this flag gives: TypeError: __init__() got an unexpected keyword argument 'store_cv_values'.

Augmented Problem Statement: Issue Description: The sklearn.linear_model.RidgeClassifierCV class is missing the store_cv_values parameter in its constructor, despite documenting the cv_values attribute that should be populated when this flag is enabled. This causes a TypeError when users attempt to use the documented functionality. Based on code analysis, the issue stems from an incomplete implementation where RidgeClassifierCV inherits from BaseRidgeCV (which likely supports store_cv_values) but doesn't expose this parameter in its own __init__ method. The documentation mentions the cv_values attribute generation depends on both store_cv_values=True and cv=None, indicating that the underlying functionality for storing cross-validation values exists but isn't properly exposed through the classifier interface. Key technical insights from analysis: RidgeCV (the regression counterpart) properly implements store_cv_values parameter support. The parameter is only compatible with cv=None (Generalized Cross-Validation / Leave-One-Out CV). There appears to be architectural inconsistency between sibling classes in the ridge module. The parent class BaseRidgeCV likely contains the shared CV value storage logic but RidgeClassifierCV doesn't properly expose or pass through the parameter. Reproduction Steps: Import numpy as np, from sklearn import linear_model as lm, n = 100, x = np.random.randn(30), y = np.random.normal(size=n), # Attempt to use store_cv_values parameter fit = lm.RidgeClassifierCV(alpha=np.arange(0.1, 1000.0, 1), normalize=True, store_cv_values=True)(x, y), Error encountered: TypeError: __init__() got an unexpected keyword argument 'store_cv_values'. Internally what happens: The RidgeClassifierCV.__init__ method signature doesn't accept store_cv_values parameter. Python class TypeError before any fitting logic can execute. Despite this, the documentation references cv_values attribute that depends on this flag. The parent class BaseRidgeCV may already have the underlying implementation but the parameter isn't properly passed through. Additional reproduction considerations: Test with proper classification target (integer or categorical labels) instead of continuous 'y'. Test with cv=None (default) to ensure compatibility with GCV mode where cv_values_storage is meaningful. Confirm behavior with RidgeCV which properly supports this parameter. Expected Behavior: RidgeClassifierCV should accept the store_cv_values boolean parameter in its constructor with default value False. When store_cv_values=True and cv=None (default GCV) the fitted object should have a populated cv_values attribute containing cross-validation values for each sample and alpha. The cv_values attribute should have shape (n_samples, n_alpha) or (n_samples, n_responses, n_alpha) depending on the problem structure. When store_cv_values=False, the cv_values attribute should be created (to save memory). Behavior should be consistent with RidgeCV which already implements this functionality correctly. Exploration Hints: Files to Examine: sklearn/linear_model/ridge.py: Why relevant: This file contains the RidgeClassifierCV class definition and is the primary location of the bug. The class inherits from BaseRidgeCV but doesn't expose the store_cv_values parameter in its constructor despite documenting the cv_values attribute. "Key findings": The RidgeCV class in the same file properly implements store_cv_values support. There's likely inheritance-related parameter exposure issues. The cv_values documentation exists but the existing parameter doesn't. sklearn/linear_model/feature_selection.py: Why relevant: Contains existing test coverage for store_cv_values functionality in RidgeCV (test_ridgecv_store_cv_values) but likely lacks equivalent tests for RidgeClassifierCV, indicating incomplete implementation coverage. "Key findings": Demonstrates the expected behavior pattern for store_cv_values. Shows dependency on cv=None for cv_values storage. Reveals API inconsistency between regression and classification variants. Key Classes/Functions: RidgeClassifierCV.__init__: What to look for: Missing store_cv_values=False parameter in method signature and lack of parameter passing to parent class. "Exploration insights": The code uses boolean cv_values attribute. May need to add validation for cv=None requirement. Should error parameter handling in RidgeCV.__init__(). BaseRidgeCV.fit: What to look for: Underlying CV value storage logic that should be triggered when store_cv_values=True. "Exploration insights": Likely contains the actual implementation that RidgeClassifierCV should leverage. May conditionally create cv_values attribute based on parameters. The delegation pattern from RidgeClassifierCV.fit to this method needs proper parameter routing. Areas of Interest: Parameter inheritance and delegation pattern: The way RidgeClassifierCV inherits classification and delegates to parent CV logic may not properly pass through store_cv_values parameter even if added to constructor. Cross-validation value storage logic: The specific conditions under which cv_values attribute gets populated (requires both store_cv_values=True AND cv=None) and how this interacts with the GCV (Generalized Cross-Validation) computation path. API consistency maintenance: Ensuring that RidgeClassifierCV exposes the same interface as RidgeCV for shared functionality to maintain user expectations and documentation consistency. Fix Hints: High Confidence Locations: sklearn/linear_model/ridge.py: RidgeClassifierCV.__init__ method: Why likely: This is the exact location where the TypeError originates. The method signature is missing the store_cv_values parameter that users expect to be able to pass. The fix requires adding the parameter to the signature and passing it to the parent constructor. Implementation Hints: Primary fix approach: 1. Add store_cv_values=False parameter to RidgeClassifierCV.__init__ signature. 2. Pass this parameter through to the parent BaseRidgeCV.__init__ call. 3. Add validation to ensure store_cv_values only works with cv=None as documented. 4. Add appropriate docstring documentation for the new parameter. "Limitations": May require ensuring that the underlying parent class properly handles the parameter for classification case, as the transformation from classification to regression problem might affect cv_values storage. Alternative approach: If the parent class doesn't properly handle store_cv_values for the classification case, the cv_values storage logic might need to be partially reimplemented or better integrated in RidgeClassifierCV.fit(). Alternative Hypotheses: Hypothesis 1: Intentional design decision with documentation error: "Reasoning": The parameter might have been intentionally excluded from RidgeClassifierCV due to theoretical or practical issues with storing CV values in classification contexts, but the documentation was not updated to reflect this. This would be due to method overloading (without proper explicit) parameter forwarding. "Evidence consideration": The documentation explicitly mentions the cv_values attribute for RidgeClassifierCV, contradicting intentional exclusion. Hypothesis 2: Inheritance chain issue: "Reasoning": RidgeClassifierCV might inherit from a different base class or use a different initialization pattern that doesn't properly expose all parent parameters, even though BaseRidgeCV (or equivalent) implements the functionality. This could be due to method overloading (without proper explicit) parameter forwarding. "Evidence consideration": Code analysis shows the fit method properly delegates to parent, suggesting initialization/parameter passing is the issue rather than execution logic.

Figure 16: CodeScout visualization of scikit-learn__scikit-learn-10297 with Qwen3-Coder.

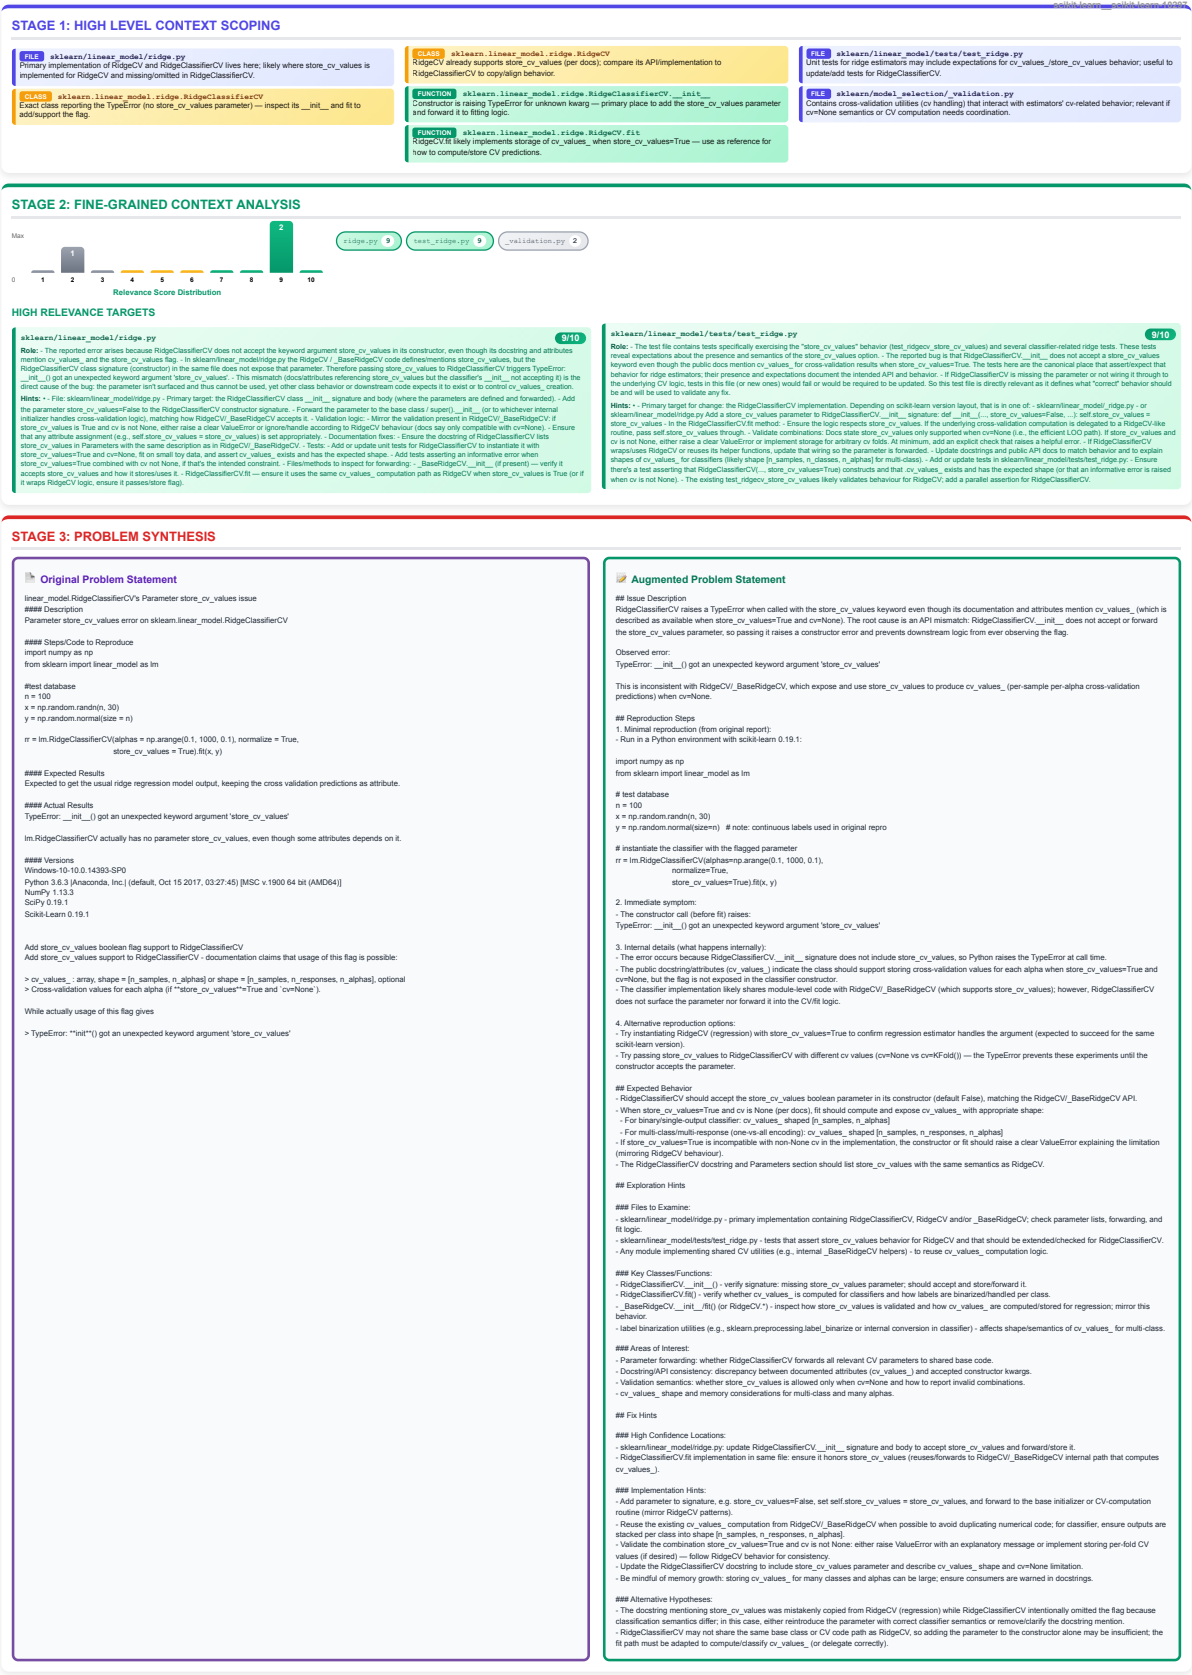


Figure 17: CodeScout visualization for scikit-learn__scikit-learn-10297 with GPT-5-mini.

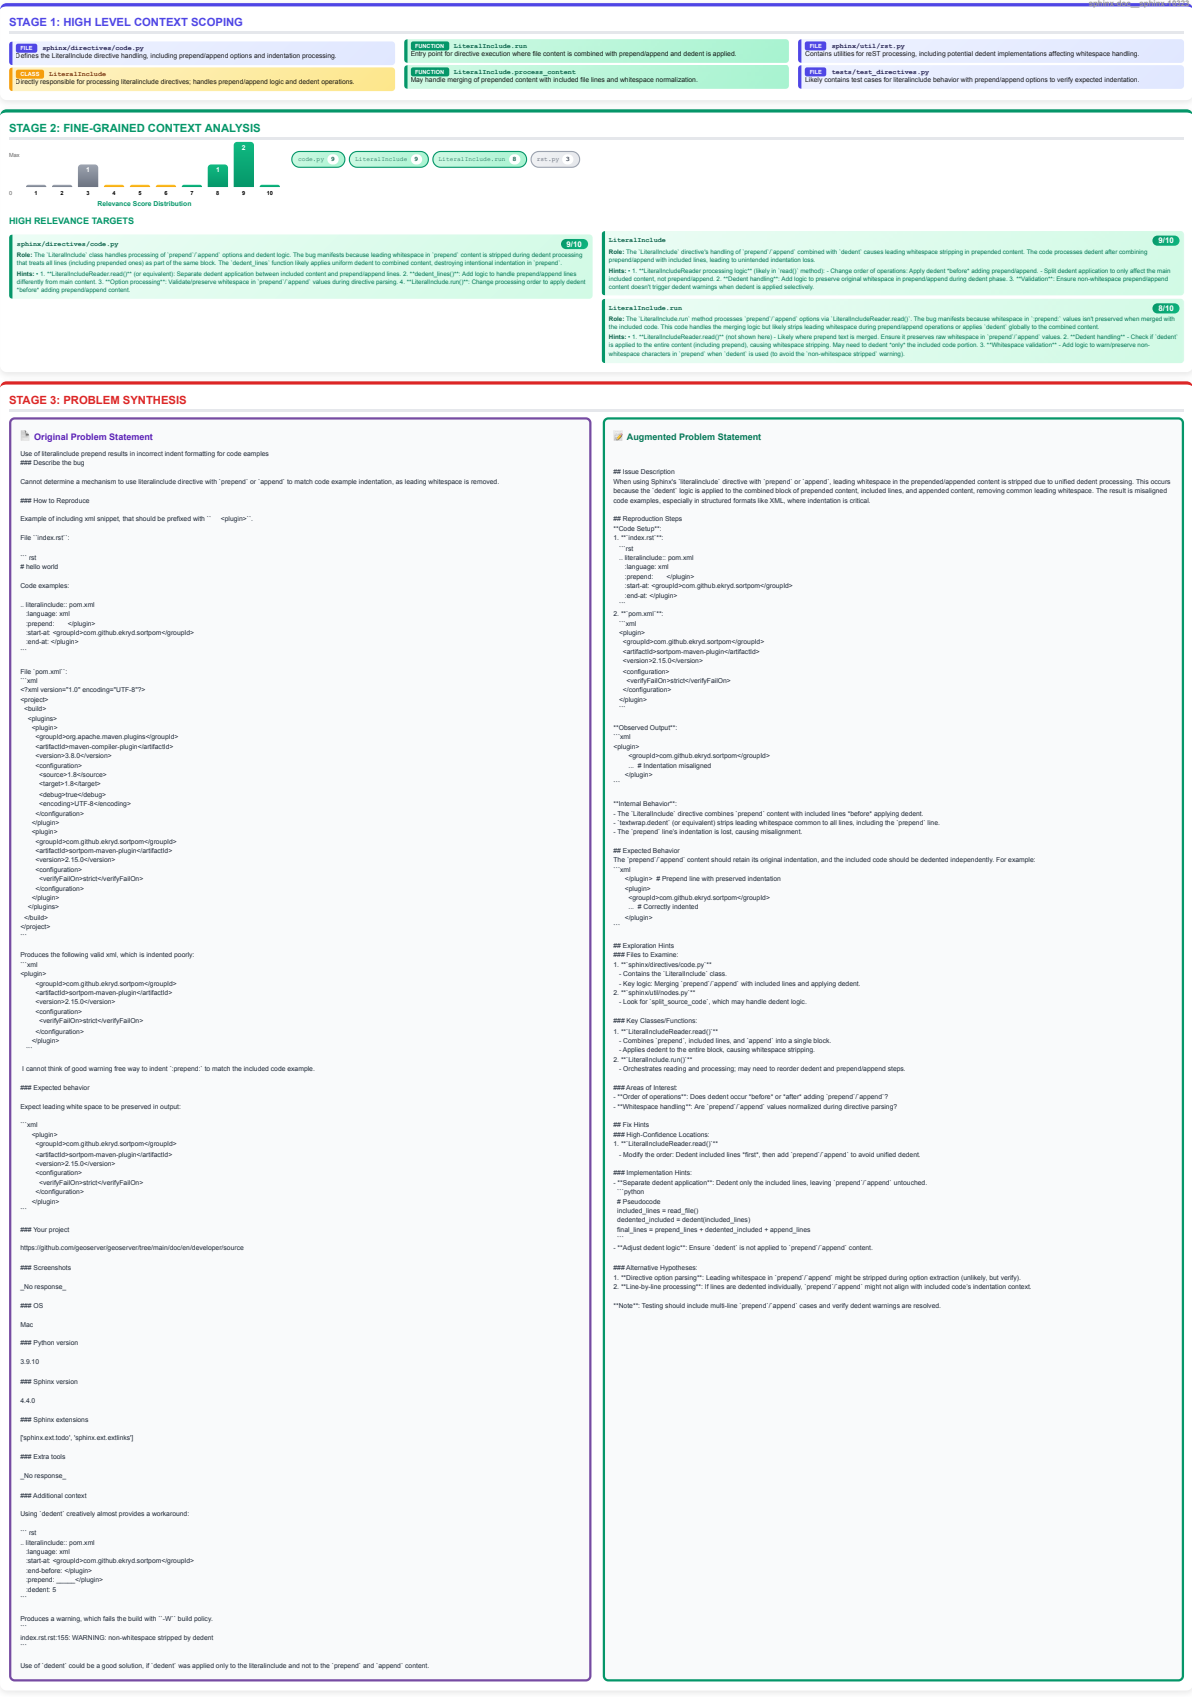


Figure 18: CodeScout visualization for sphinx-doc__sphinx-10323 with DeepSeek-R1.

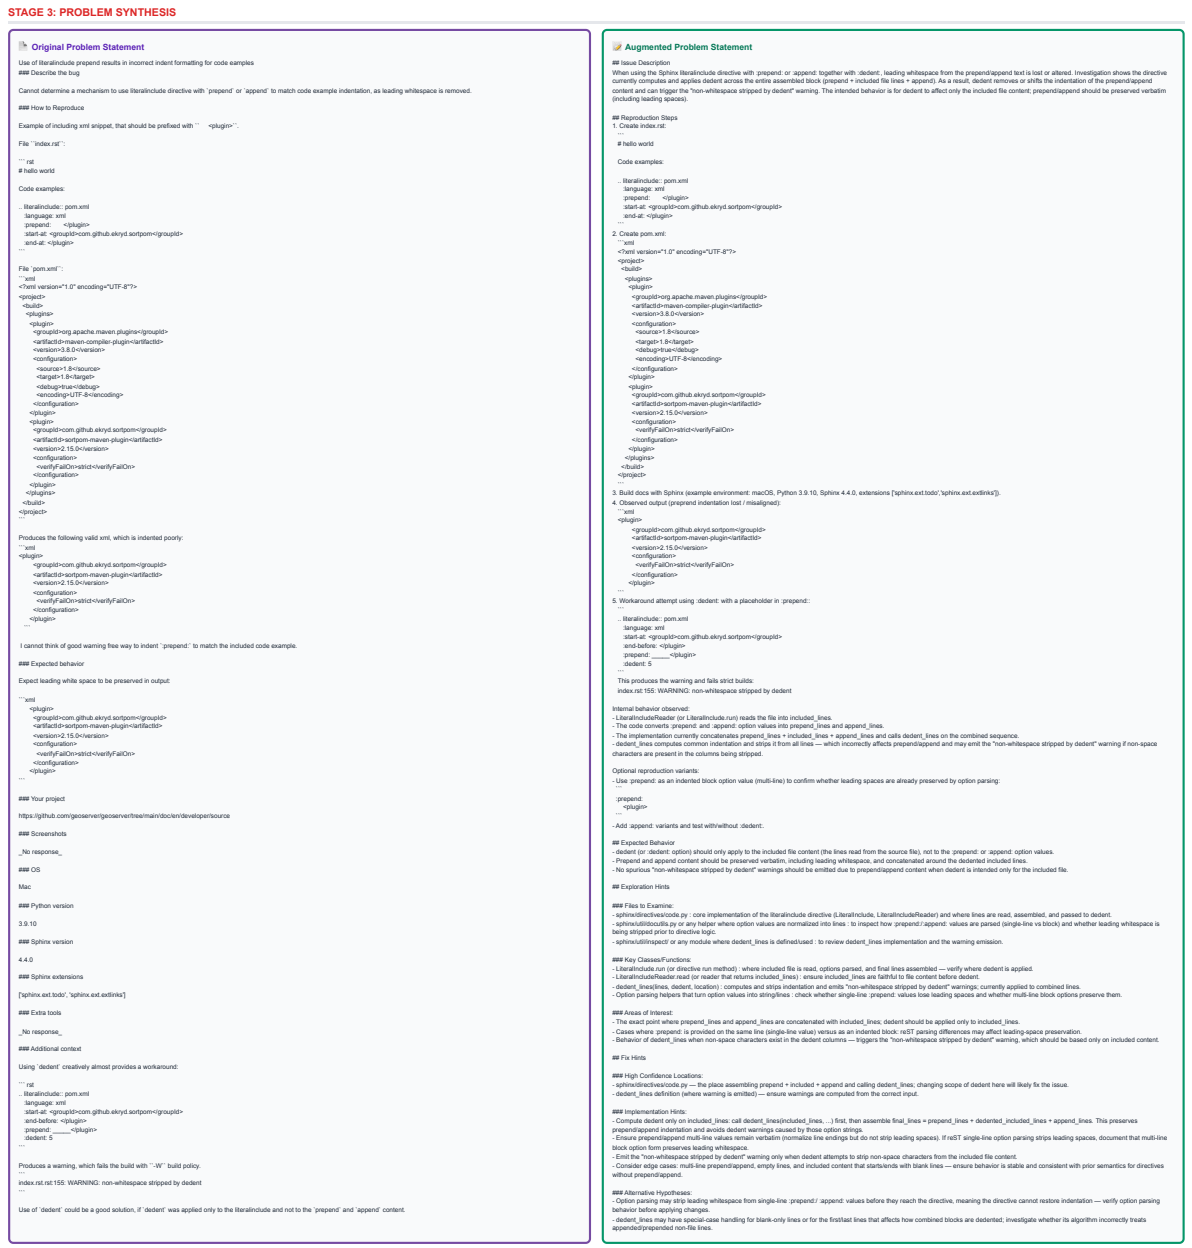
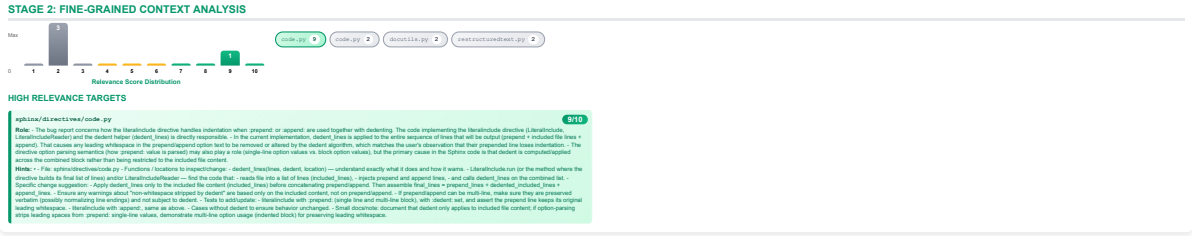
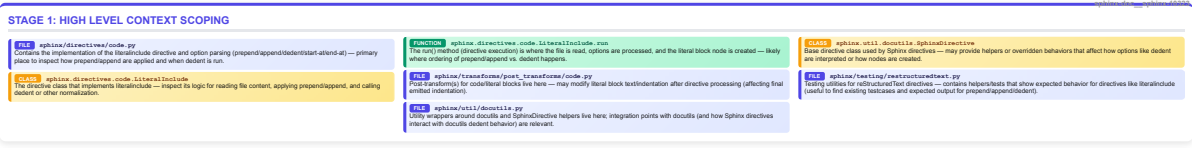


Figure 20: CodeScout visualization for sphinx-doc__sphinx-10323 with GPT-5-mini.

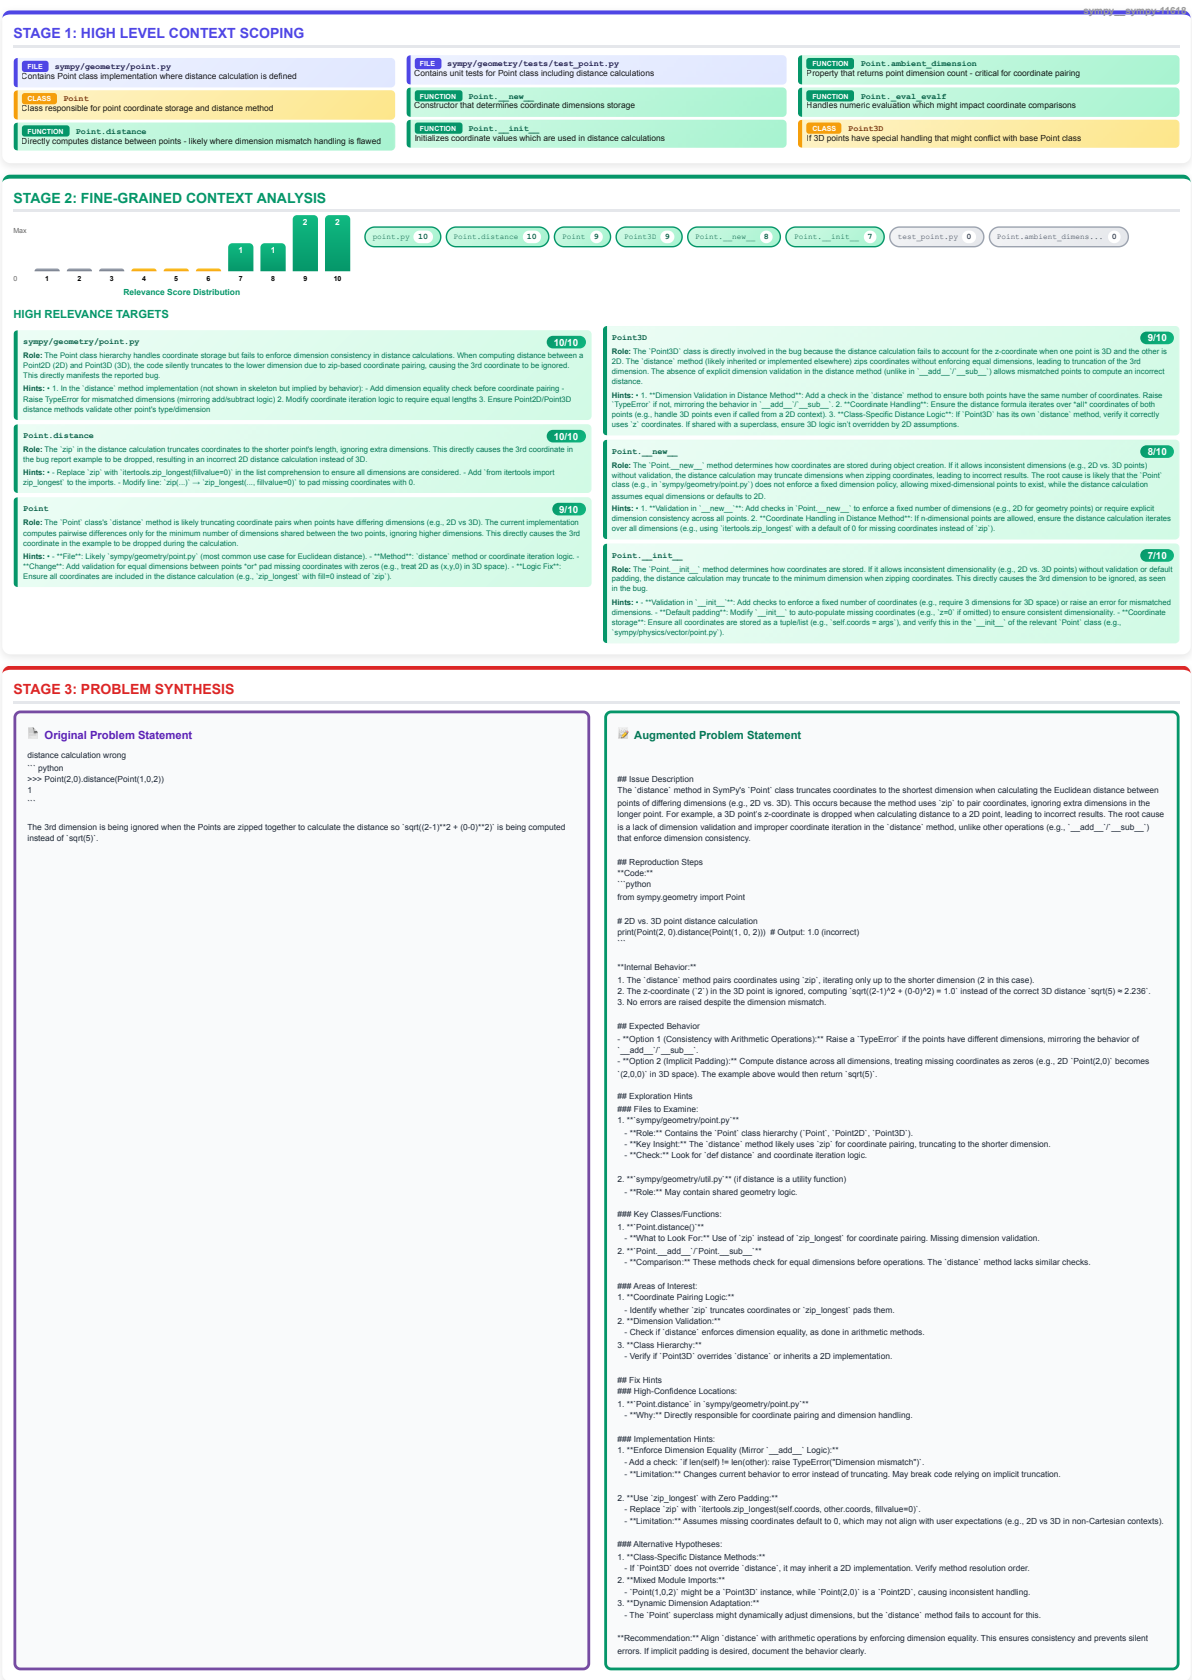


Figure 21: CodeScout visualization for sympy__sympy-11618 with DeepSeek-R1.

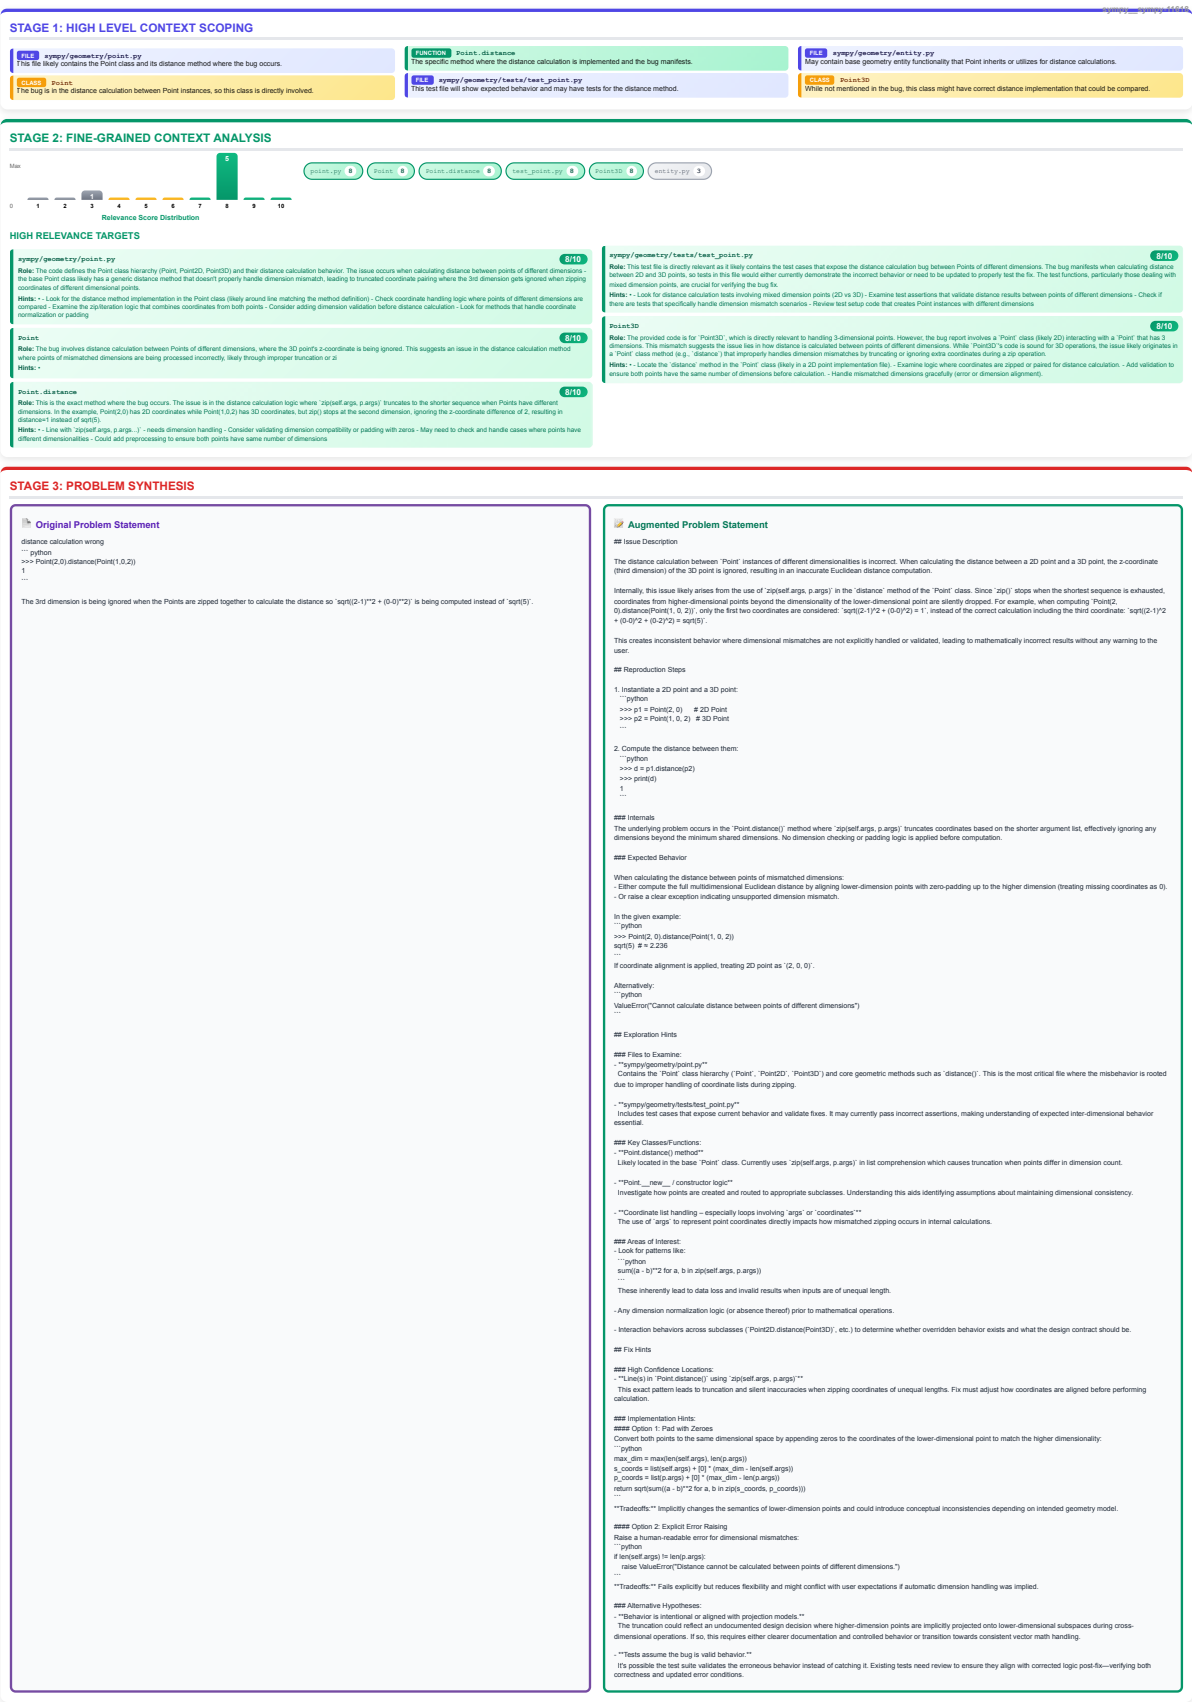


Figure 22: CodeScout visualization for sympy__sympy-11618 with Qwen3-Coder.

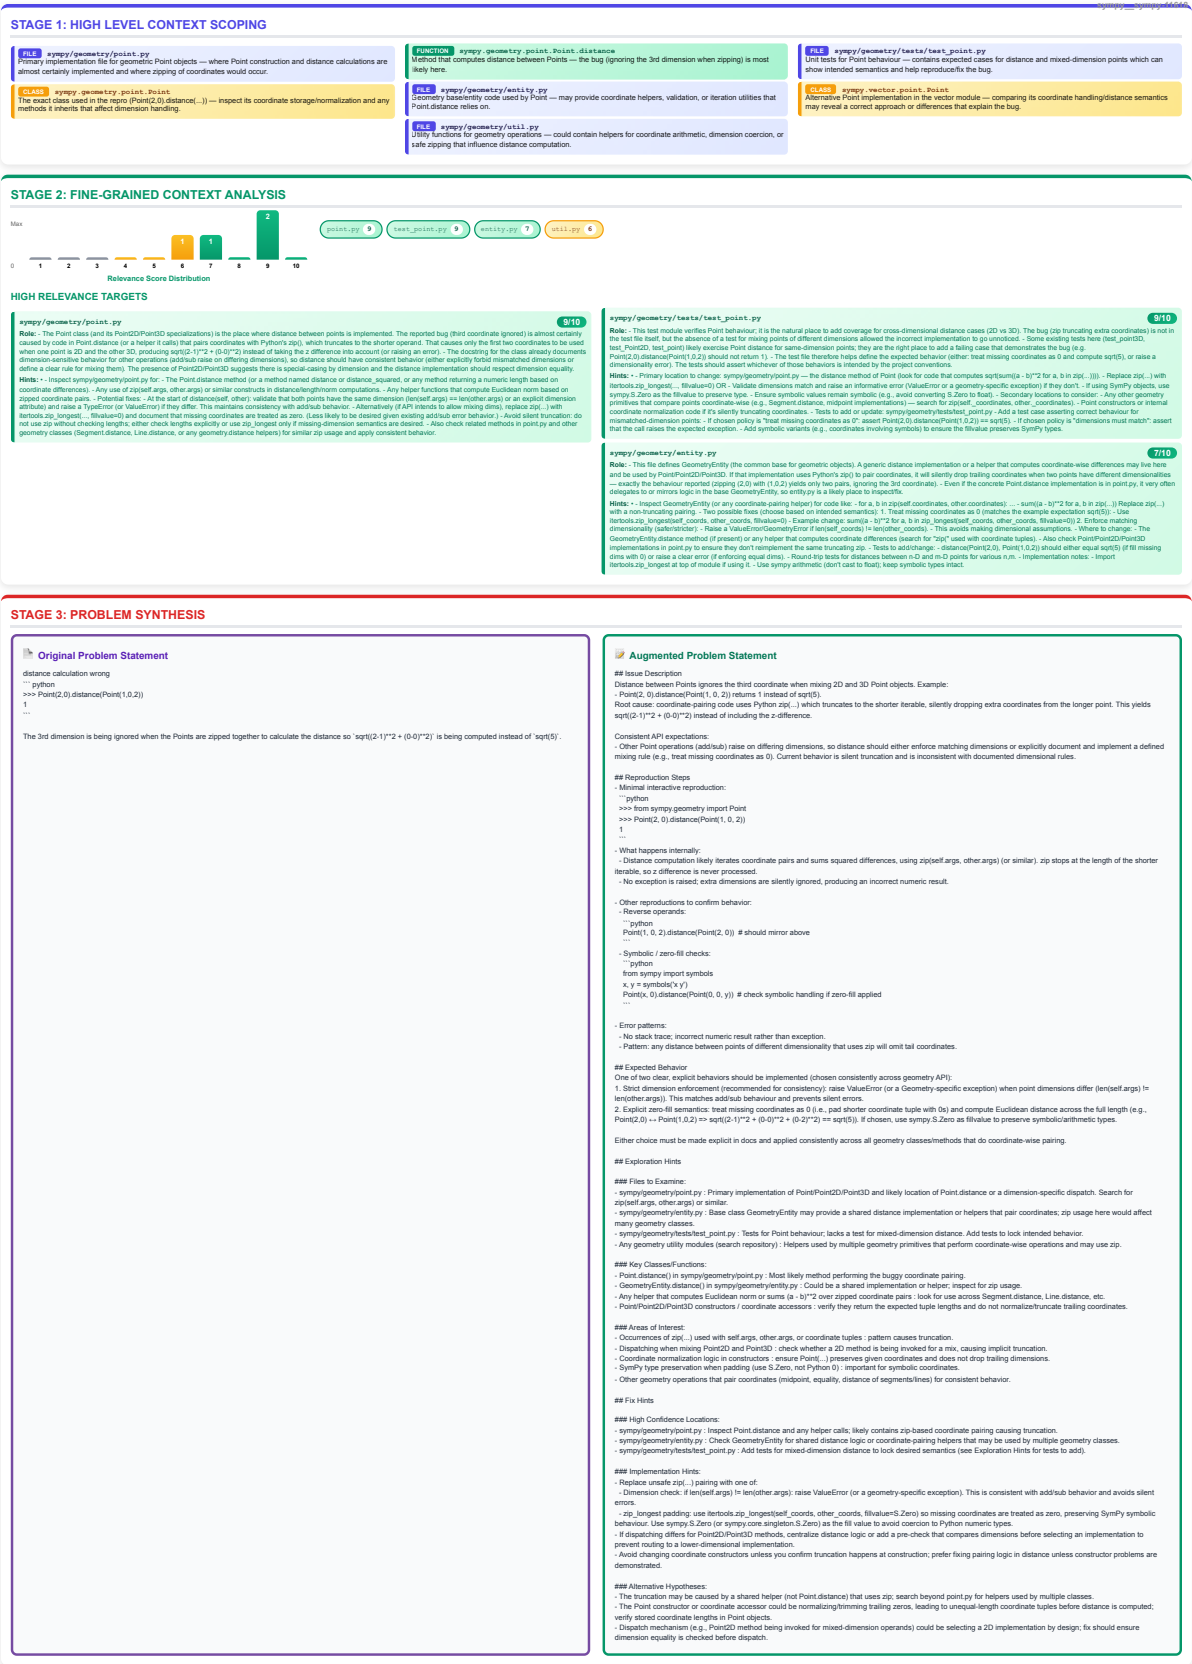


Figure 23: CodeScout visualization for sympy__sympy-11618 with GPT-5-mini.

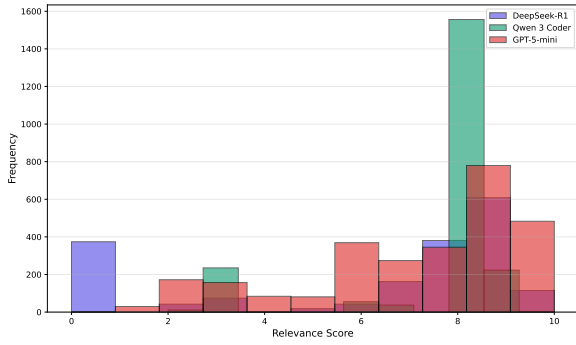


Figure 24: CodeScout relevance score distribution across all three methods.

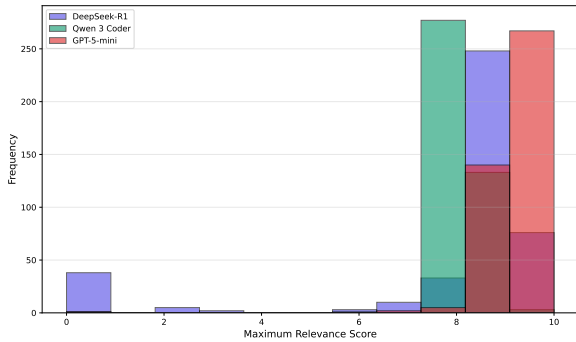


Figure 25: CodeScout maximum relevance score distribution per instance.

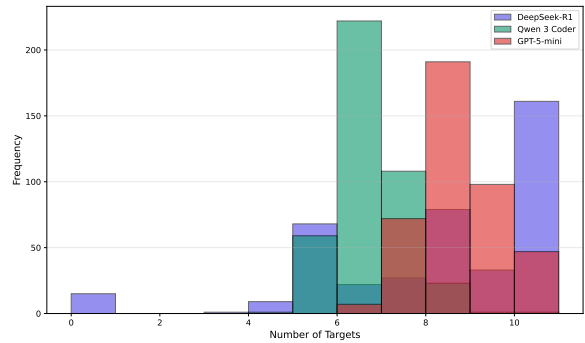


Figure 26: CodeScout number of targets distribution per instance.

the overlap of unique targets identified across methods, showing that a substantial portion of targets are method-specific, with GPT-5-mini identifying the largest unique set.

Agreement Analysis. Figure 28 presents Bland-Altman plots comparing the average scores between method pairs. These plots reveal systematic differences: DeepSeek-R1 consistently scores lower than both Qwen3-Coder and GPT-5-mini, as evidenced by the negative mean differences. The agreement limits show the range of score differences across instances.

Score Correlations. Figure 29 shows heatmaps of score agreement for targets that were analyzed by multiple methods. Each cell shows the count of targets where the row method assigned one score and the column method assigned another. The diagonal entries represent exact agreement, while off-diagonal entries indicate disagreement. The conditional probabilities in the colorbar show how likely one method is to assign a particular score given the other method's score.

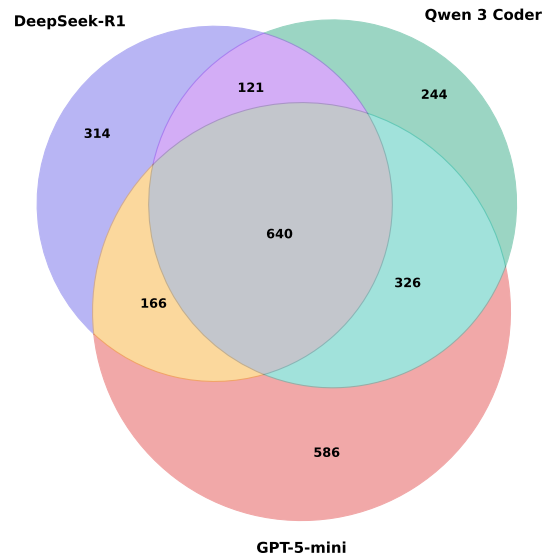
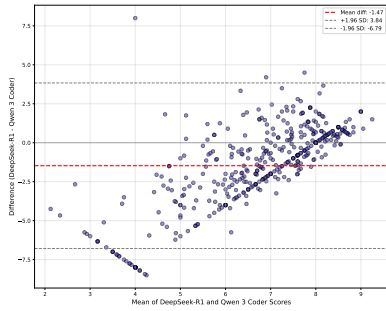
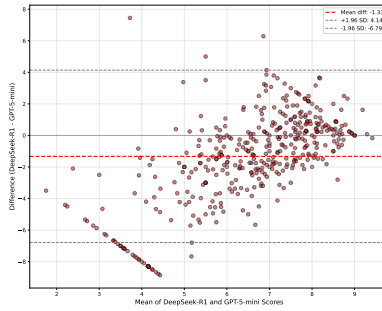


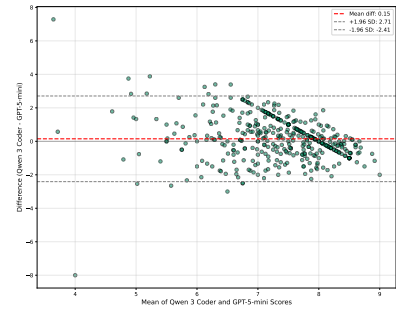
Figure 27: CodeScout target coverage: Venn diagram showing unique and shared exploration targets.



(a) DeepSeek-R1 vs Qwen3-Coder

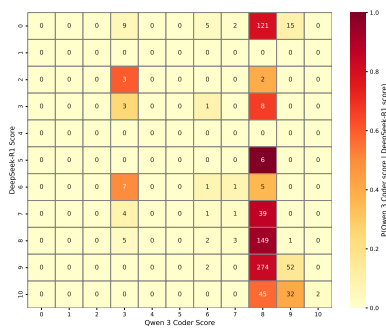


(b) DeepSeek-R1 vs GPT-5-mini

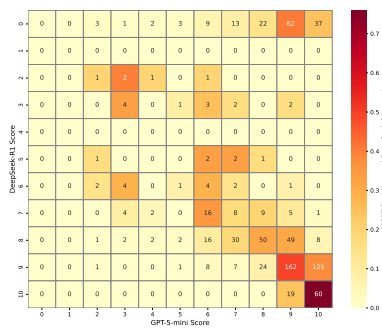


(c) Qwen3-Coder vs GPT-5-mini

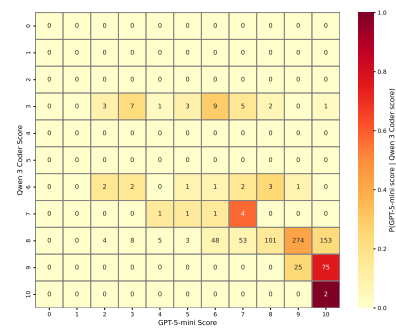
Figure 28: CodeScout agreement analysis: Bland-Altman plots comparing average scores between LLMs.



(a) DeepSeek-R1 vs Qwen3-Coder



(b) DeepSeek-R1 vs GPT-5-mini



(c) Qwen3-Coder vs GPT-5-mini

Figure 29: CodeScout score agreement heatmaps between methods for common targets.