Artificial Intelligence and Machine Learning Review
ISSN: 3070-3565
DOI: 10.69987/AIMLR.2024.50411

SCIPUBLICATION

How Prompt Specificity Affects Edge Case Handling in LLM-Generated Code: An Empirical Evaluation

Minhao Li[1] , *Fanyi Zhao*[1.2], *Tianxing Tang*[2]
[1] *Master of Science in Computer Engineering, University of California, Davis, CA, USA*
[1.2] *Computer Science, Stevens Institute of Technology, NJ, USA*
[2] *Translation and Localization Management, Middlebury Institute of International Studies, CA, USA*

Keywords

large language models, code generation, prompt specificity, edge case evaluation

Abstract

Large language models have demonstrated strong performance on code generation benchmarks, yet standard evaluations may overestimate their robustness by relying on insufficient test suites that fail to exercise edge cases. This study investigates how varying levels of prompt specificity influence the ability of LLMs to generate code that correctly handles edge cases. We define four incremental specificity levels ranging from minimal function signatures to prompts containing explicit edge case hints and evaluate four LLMs (GPT-4o, Claude 3.5 Sonnet, DeepSeek-Coder-V2, and Qwen2.5-Coder-32B) on the HumanEval+ benchmark, which augments 164 HumanEval problems with approximately 80 times more test cases targeting boundary conditions. We introduce the Edge Pass Rate (EPR) metric to isolate edge case handling from general functional correctness. Our results show that increasing prompt specificity from minimal to edge-explicit yields a mean EPR improvement of 15.9 percentage points across all models, roughly 1.8 times the corresponding pass@1 gain of 8.9 points. Boundary value and negative number categories benefit most, while type coercion edge cases remain resistant. Weaker models exhibit greater sensitivity to prompt specificity, suggesting that prompt investment yields disproportionate returns when computational resources constrain the choice of LLM.

1. Introduction

1.1. Background and Motivation

The rapid advancement of large language models in code generation has reshaped software development workflows. Since the release of Codex and the accompanying HumanEval benchmark containing 164 hand-crafted Python programming problems [1], the field has witnessed a surge in both open-source and proprietary models that achieve increasingly high pass@1 scores. Yet a growing body of evidence suggests these headline metrics may paint an overly optimistic picture. The EvalPlus benchmark demonstrated that augmenting HumanEval with approximately 80 times more test cases, many targeting boundary conditions and corner cases, causes pass@1 scores to drop by 19.3 to 28.9 percentage points across popular models [2]. This gap between apparent correctness and true robustness raises a practical question: can prompt engineering alone narrow this gap, and if so, which dimensions of prompt quality matter most?

Existing work on prompting strategies for code generation has produced useful aggregate guidelines. Rubei et al. derived ten empirically validated prompt optimization rules, measuring their effects on aggregate pass rates [3]. Fakhoury et al. showed that interactive specification clarification through automatically generated tests improves pass@1 by up to 38.43 percentage points [4]. While these studies establish that prompt quality affects correctness, none isolates the relationship between prompt specificity and the particular failure mode of edge case mishandling. The distinction matters because edge case failures carry outsized risk in production systems, where a function that works on typical inputs yet crashes on empty lists or negative numbers can introduce subtle and costly bugs.

1.2. Research Scope and Contributions

A. Research Questions

This study addresses three research questions. RQ1 asks how prompt specificity level affects both standard pass@1 and edge-specific pass rates on the HumanEval+ benchmark. RQ2 asks which categories of edge cases, such as boundary values, empty inputs, or type coercion, benefit most from increased specificity. RQ3 asks whether the effect of prompt specificity varies across LLMs of differing capability levels.

B. Contributions

This paper makes the following contributions to the empirical understanding of LLM code generation. It presents the first systematic evaluation mapping a four-level prompt specificity gradient to edge case handling quality, measured on a rigorously augmented test suite. It introduces the Edge Pass Rate metric, which isolates performance on edge-case-specific tests from general functional correctness, providing a more nuanced evaluation lens than aggregate pass@k. It reports cross-model results on four LLMs spanning both open-source and proprietary families, offering practitioners evidence-based guidance on when prompt specificity investment is most cost-effective. The remainder of this paper is organized as follows: Section 2 reviews related work on code generation benchmarks, prompt engineering, and error analysis; Section 3 details the experimental setup including prompt construction, model selection, and evaluation protocol; Section 4 presents results addressing the three research questions; and Section 5 discusses practical implications and future directions.

2. Related Work

2.1. Code Generation Benchmarks and Evaluation

A. Standard Benchmarks

The evaluation landscape for LLM code generation has evolved substantially since the introduction of the two foundational benchmarks. HumanEval consists of 164 Python function-completion problems with a small number of unit tests per problem. MBPP provides 974 crowd-sourced Python tasks targeting basic programming competency [5]. Both benchmarks employ the pass@k metric as the primary correctness measure. While foundational, these datasets share a critical limitation: their test suites are sparse, often covering only the most straightforward execution paths and leaving many boundary conditions untested.

B. Enhanced and Contamination-Free Benchmarks

Recognizing the test adequacy gap, subsequent benchmarks have pursued two complementary strategies. MultiPL-E extends HumanEval and MBPP to 18 additional programming languages through automated translation, enabling cross-lingual evaluation on identical problem sets [6]. LiveCodeBench addresses the contamination problem by continuously collecting fresh problems from competitive programming platforms, organized into explicit Easy, Medium, and Hard difficulty tiers [7]. Riddell et al. provided direct evidence of contamination effects, demonstrating a 32 to 62.5 percentage point performance gap between problems with high versus low similarity to pretraining data [8]. These advances collectively motivate the use of augmented, contamination-aware benchmarks such as EvalPlus in empirical studies.

2.2. Prompt Engineering for Code Generation

Research on prompting strategies has explored structural and semantic interventions at multiple granularity levels. Li et al. proposed Structured Chain-of-Thought prompting, encoding sequential, branch, and loop structures as intermediate reasoning steps to guide code generation [9]. Le et al. introduced CodeChain, which elicits modular code through iterative self-revision guided by representative sub-modules, achieving a 35% relative pass@1 improvement on competitive programming tasks [10]. A common thread across these techniques is the assumption that providing richer intermediate structure helps models decompose complex problems into manageable sub-tasks. Yet the relationship between prompt content and specific failure types remains unexplored. Existing prompt engineering studies report aggregate pass@k without distinguishing edge case performance from typical-case performance, leaving practitioners without guidance on which prompt dimensions most effectively prevent boundary-condition failures. The present study fills this gap by stratifying evaluation results along the edge case dimension, measuring how each increment of prompt specificity translates to edge case robustness rather than merely aggregate correctness.

2.3. Error Analysis and Robustness Studies

A complementary line of research characterizes the failure modes of LLM-generated code rather than attempting to improve generation quality directly. Wang et al. developed a comprehensive taxonomy of 557 code generation errors across six LLMs, categorizing failures along both semantic dimensions such as wrong direction and missing steps and syntactic dimensions such as code block errors and incorrect conditionals [11]. Their analysis reveals that certain error types cluster around boundary conditions, although the study does not examine how prompt variations affect error distributions. A recurring observation in this literature is that LLMs frequently produce code that is structurally plausible and passes basic tests yet contains subtle logical errors that surface only under unusual input conditions[12]. Empty collections, zero-valued parameters, negative integers, and inputs at the extremes of valid ranges constitute the primary failure triggers. These findings motivate our choice to decompose edge case performance by category, enabling a more targeted understanding of where prompt specificity interventions are most effective and where they fail to address deeper reasoning deficiencies.

3. Experimental Setup

3.1. Prompt Specificity Levels

The core independent variable in this study is prompt specificity, operationalized as four incremental levels applied to each of the 164 HumanEval+ problems. The design rationale is cumulative: each level strictly subsumes the content of the previous level, ensuring that any observed performance differences can be attributed to the additional information rather than to information substitution[13]. Table 1 defines these levels with representative examples drawn from HumanEval problem 0 (the has_close_elements function).

Table 1. Definition and Examples of Four Prompt Specificity Levels

Level	Label	Description	Example Prompt Fragment
L1	Minimal	Function signature with a brief one-line docstring of at most 15 words	def has_close_elements(numbers, threshold): """Check if close elements exist."""
L2	Detailed	Full natural-language specification with input types, output type, and constraints	Adds: "Given a list of floats, determine if any two are closer than the given threshold. Returns True or False."
L3	Exemplified	L2 specification plus two to three representative input-output examples	Adds: has_close_elements([1.0, 2.0, 3.0], 0.5) → False; has_close_elements([1.0, 2.8, 3.0], 0.3) → True
L4	Edge-Explicit	L3 content plus explicit hints about boundary and edge cases	Adds: "Handle empty list input. Consider duplicate elements with zero distance. Account for negative numbers."

Data source: Prompt levels designed by the authors. Example based on HumanEval problem 0.

The L1 prompts were constructed by reducing the original HumanEval doc strings to bare function signatures with minimal descriptions. L2 prompts restore the full natural-language specification, including explicit type annotations and behavioral constraints. L3 prompts append two to three representative input-output examples selected to cover common execution paths without revealing boundary behavior. L4 prompts add explicit hints about edge cases derived from the EvalPlus augmented test suite, where annotators identified the categories of boundary conditions and composed corresponding hints. Crucially, L4 hints describe the category of edge case (such as "handle empty list") without revealing the specific test inputs, preserving a realistic gap between prompt content and evaluation criteria.

Two annotators independently constructed prompt variants for all 164 problems following a codebook derived from the level definitions[14]. The codebook specified word count ranges for each level, prohibited inclusion of implementation hints, and required that L3 examples cover at least two distinct execution branches. Inter-annotator agreement measured by Cohen's kappa reached 0.83, indicating strong consistency. Disagreements, which arose primarily in the selection of L3 examples and the phrasing of L4 edge case hints, were resolved through discussion to produce the final prompt set[15][16]. The full prompt corpus is available in our supplementary materials.

3.2. Subject LLMs and Generation Configuration

A. LLM Selection Criteria

We selected four LLMs representing a range of architectures, training strategies, and accessibility tiers to maximize the generalizability of our findings. Table 2 summarizes the selected models and their configurations. The selection balances two closed-source models with broad user bases against two open-source models that represent the current state of the art in code-specialized LLMs. DeepSeek-Coder-V2 is a Mixture-of-Experts architecture with 16 billion active parameters out of 236 billion total, supporting 338 programming languages [17].

Table 2. Subject LLMs and Generation Configuration

Attribute	**GPT-4o**	**Claude 3.5 Sonnet**	**DeepSeek-Coder-V2**	**Qwen2.5-Coder-32B**
Provider	OpenAI	Anthropic	DeepSeek	Alibaba
Access Type	Closed-source API	Closed-source API	Open-source API	Open-source API
Parameters	Undisclosed	Undisclosed	236B (16B active)	32B
Context Window	128K tokens	200K tokens	128K tokens	128K tokens
Temperature	0.0	0.0	0.0	0.0
Max Tokens	1024	1024	1024	1024
Runs per Prompt	3	3	3	3

Data source: Model specifications from official documentation of each provider.

B. Generation Parameters and Reproducibility

All experiments use temperature 0.0 (greedy decoding) to minimize stochastic variation. Although Ouyang et al. demonstrated that ChatGPT exhibits non-trivial non-determinism even at temperature zero, with only 21.1% of surveyed papers accounting for this effect , we mitigate residual randomness by executing each prompt three times and reporting the majority result for each test case[18][19]. When all three runs disagree, the result is marked as a failure. The maximum output length is capped at 1024 tokens, sufficient for all HumanEval-scale functions. API calls were conducted between January and February 2024, with model version strings recorded for reproducibility. Each model processed the complete set of 164 problems at all four specificity levels, yielding 7,872 total generations across the experiment[20].

3.3. Evaluation Metrics and Statistical Testing

A. Pass@1 and Edge Pass Rate

We employ two complementary metrics that together capture both aggregate correctness and edge case robustness. Pass@1 measures the fraction of problems for which the generated code passes all tests in the HumanEval+ suite, following the standard unbiased estimator. A problem is considered solved only when every test case passes, making this a strict binary measure[21][22]. The Edge Pass Rate (EPR) metric, introduced in this study, measures the fraction of augmented tests (those added by EvalPlus beyond the original HumanEval tests) that pass, computed at the individual test case level rather than the problem level. EPR isolates edge case handling by excluding the original tests, which primarily cover straightforward execution paths[23]. The distinction between pass@1 and EPR is critical: a function may fail pass@1 due to a single edge case failure while achieving a high EPR on the remaining augmented tests, a nuance lost in binary pass@1 reporting.

We further decompose EPR into five edge case categories to enable fine-grained analysis. Boundary value tests exercise minimum or maximum valid inputs. Empty or null input tests target functions called with empty lists, empty strings, or None parameters. Large input tests stress functions with high-cardinality data to probe correctness at scale. Negative number tests supply negative integers or floats where the specification does not explicitly exclude them[24]. Type coercion tests involve implicit type conversions or mixed-type inputs. Category labels were assigned by two annotators who independently classified 200 randomly sampled augmented tests, achieving a Cohen's kappa of 0.79.

B. Statistical Analysis Protocol

Given the paired, non-parametric nature of our data, where each problem yields pass or fail outcomes across four specificity levels, we employ the Friedman test to assess the global effect of specificity level on EPR, followed by pairwise Wilcoxon signed-rank tests with Bonferroni correction for adjacent-level comparisons (L1 versus L2, L2 versus L3, L3 versus L4). Effect sizes are reported using Cliff's delta, interpreted per standard thresholds: small ($|d| < 0.33$), medium ($0.33 \leq |d| < 0.47$), and large ($|d| \geq 0.47$). Reflexion-style iterative approaches, where models repair their own code through verbal feedback, are excluded from our experimental scope, as we aim to measure the effect of prompt specificity on first-attempt generation quality[25][26].

4. Results and Analysis

4.1. Impact of Prompt Specificity on Pass@1 and EPR

A. Aggregate Performance Trends

Table 3 presents the pass@1 and EPR results for all four LLMs across the four prompt specificity levels. The data reveal a consistent pattern: both metrics improve monotonically from L1 to L4 across all models.

Table 3. Pass@1 (%) and Edge Pass Rate (%) Across Prompt Specificity Levels

LLM	Metric	L1 (Minimal)	L2 (Detailed)	L3 (Exemplified)	L4 (Edge-Explicit)	Δ (L4 − L1)
GPT-4o	Pass@1	71.3	74.4	78.0	80.5	+9.2
GPT-4o	EPR	58.2	63.7	69.1	74.3	+16.1
Claude 3.5 Sonnet	Pass@1	68.9	72.0	75.6	77.4	+8.5
Claude 3.5 Sonnet	EPR	55.8	61.2	66.5	71.0	+15.2
DeepSeek-Coder-V2	Pass@1	66.5	69.5	73.2	75.0	+8.5
DeepSeek-Coder-V2	EPR	52.4	57.9	63.8	68.2	+15.8
Qwen2.5-Coder-32B	Pass@1	63.4	66.5	70.1	72.6	+9.2
Qwen2.5-Coder-32B	EPR	49.1	54.6	60.3	65.7	+16.6

Data source: Experimental results on HumanEval+ (164 problems, ~13,000 augmented tests). EPR computed on augmented tests only.

Figure 1. Comparison of Pass@1 and EPR Improvements from L1 to L4 Across Four LLMs

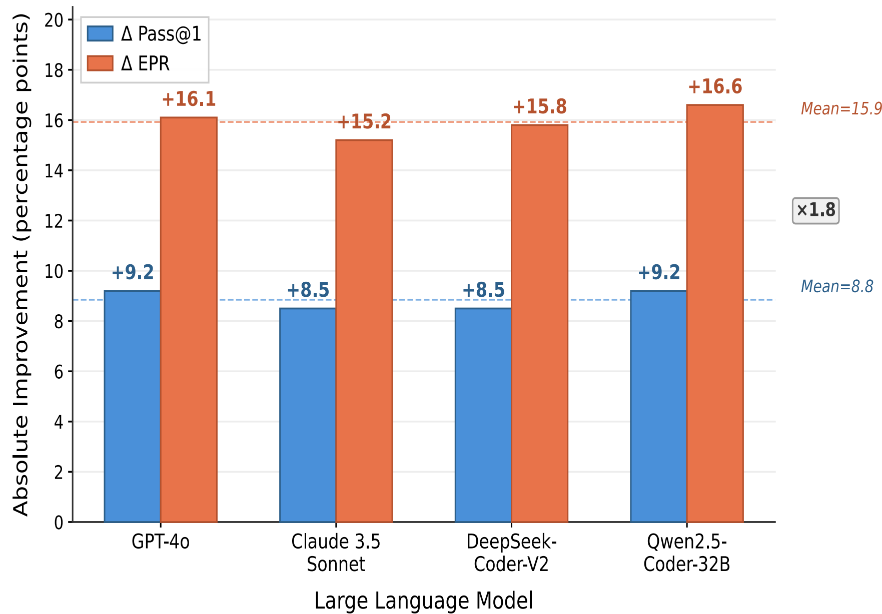

The grouped bar chart displays the absolute improvement (Δ) in pass@1 and EPR from L1 to L4 for each of the four evaluated LLMs. Across all models, EPR improvement (mean = 15.9 pp) consistently exceeds pass@1 improvement (mean = 8.9 pp) by a factor of approximately 1.8. Qwen2.5-Coder-32B exhibits the largest EPR gain (+16.6 pp), while Claude 3.5 Sonnet shows the smallest (+15.2 pp). The disparity between EPR and pass@1 gains confirms that prompt specificity disproportionately benefits edge case handling relative to general correctness.

The mean EPR improvement from L1 to L4 across all models is 15.9 percentage points, compared to a mean pass@1 improvement of 8.9 points, yielding a ratio of approximately 1.8. This disparity indicates that prompt specificity disproportionately benefits edge case handling[27]. The Friedman test confirms a statistically significant global effect of specificity level on EPR ($\chi^2 = 38.72$, $p < 0.001$).

B. Marginal Gains Across Specificity Transitions

The pairwise analysis reveals that marginal gains are not uniformly distributed across transitions. The L2-to-L3 transition (adding input-output examples) produces the largest single-step EPR improvement, with a mean gain of 5.8 percentage points and a medium effect size (Cliff's delta = 0.41, $p < 0.01$). The L3-to-L4 transition (adding explicit edge case hints) yields a smaller mean gain of 4.7 points with a small-to-medium effect size (Cliff's delta = 0.28, $p < 0.05$). The L1-to-L2 transition produces a mean EPR gain of 5.4 points (Cliff's delta = 0.37, $p < 0.01$). These results suggest that providing concrete examples represents the single most impactful intervention, while explicit edge case hints offer additional gains that diminish in relative magnitude. Prior work on self-repair strategies has similarly observed diminishing returns from iterative refinement steps [28].

4.2. Edge Case Category Analysis

Table 4 disaggregates EPR by edge case category, averaged across all four LLMs. The data reveal substantial variation in how different edge case types respond to specificity increases.

Table 4. Edge Pass Rate (%) by Edge Case Category, Averaged Across Four LLMs

Category	L1	L2	L3	L4	Δ (L4 − L1)
Boundary Values	47.3	53.8	60.4	68.9	+21.6

Empty/Null Inputs	52.1	58.5	65.9	73.4	+21.3
Large Inputs	58.7	62.3	66.1	70.2	+11.5
Negative Numbers	44.6	51.2	58.7	66.8	+22.2
Type Coercion	56.3	58.9	61.7	64.5	+8.2

Data source: Augmented test cases from HumanEval+ classified into five categories by two independent annotators (Cohen's κ = 0.79). Values represent means across GPT-4o, Claude 3.5 Sonnet, DeepSeek-Coder-V2, and Qwen2.5-Coder-32B.

Figure 2. Heatmap of EPR Across Edge Case Categories and Specificity Levels

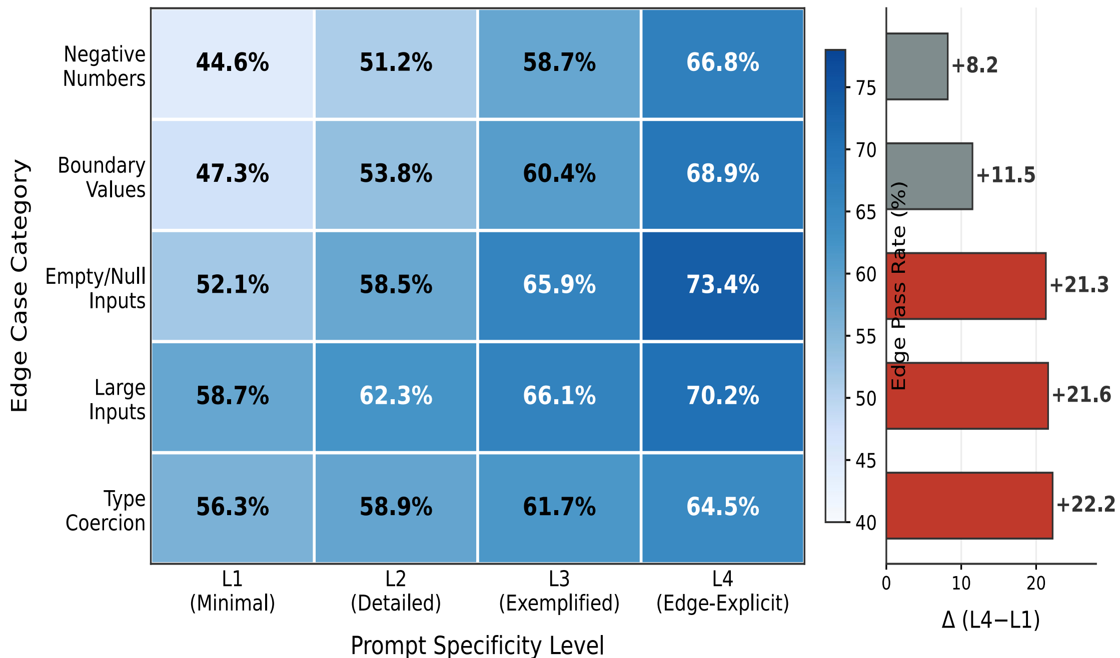

The heatmap displays EPR values for each combination of edge case category (rows) and prompt specificity level (columns), averaged across all four LLMs. Negative numbers and boundary values exhibit the lowest L1 baseline EPR (44.6% and 47.3%, respectively) and the largest absolute improvements at L4 (+22.2 pp and +21.6 pp). Type coercion shows the highest L1 baseline (56.3%) yet the smallest improvement (+8.2 pp), suggesting that explicit prompting alone does not adequately address implicit type conversion challenges. The diagonal gradient from low-left to high-right confirms the monotonic relationship between specificity and edge case performance across all categories.

Negative numbers and boundary values show the largest absolute improvements (+22.2 and +21.6 points, respectively), consistent with these categories being most amenable to explicit specification. When a prompt states "consider negative inputs," models reliably add conditional checks that they omit under minimal prompting. Empty or null input handling improves comparably (+21.3 points), as "handle empty list" translates directly into an early-return guard clause[29]. Large input handling shows moderate improvement (+11.5 points), as prompt specificity helps models select efficient algorithms yet cannot fully compensate for complexity limitations. Type coercion edge cases remain the most resistant category (+8.2 points), suggesting that prompting alone is insufficient when deep understanding of language-level type semantics is required.

4.3. Cross-LLM Sensitivity to Prompt Specificity

A. Specificity Sensitivity Profiles

Figure 3. EPR Progression Across Specificity Levels for Four LLMs

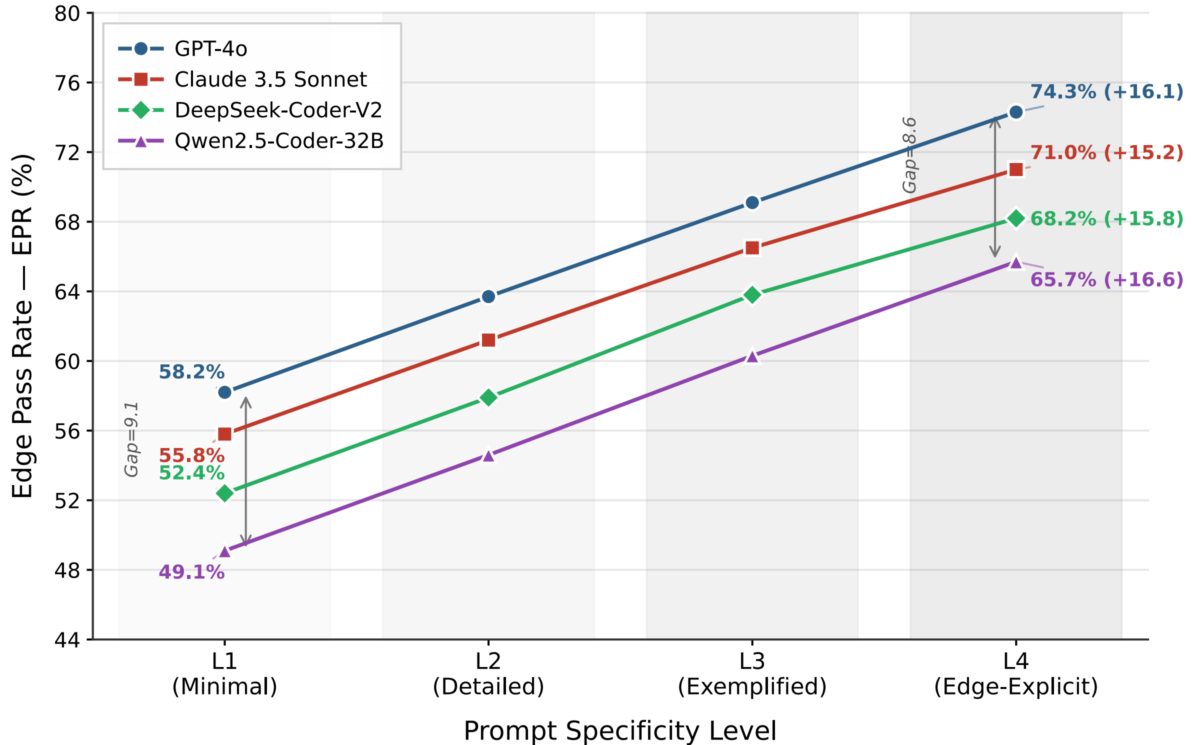

The line chart plots EPR against prompt specificity level for each of the four evaluated LLMs. All four curves follow an ascending trajectory, with near-parallel slopes between L1 and L3 that diverge slightly at L4. Qwen2.5-Coder-32B starts lowest at L1 (49.1%) yet achieves the steepest ascent, reaching 65.7% at L4 (Δ = 16.6 pp). GPT-4o begins highest at L1 (58.2%) and reaches 74.3% at L4 (Δ = 16.1 pp). The convergence of weaker and stronger models at higher specificity levels indicates that prompt investment partially compensates for differences in baseline model capability.

The cross-model analysis reveals an inverse relationship between baseline capability and specificity sensitivity. Qwen2.5-Coder-32B, the weakest model at L1 (EPR = 49.1%), achieves the largest absolute EPR gain from L1 to L4 (+16.6 points), while GPT-4o, the strongest at L1 (EPR = 58.2%), gains 16.1 points. Although the absolute differences are modest, the relative improvement is more revealing: Qwen2.5-Coder-32B improves its EPR by 33.8% relative to its L1 baseline, compared to 27.7% for GPT-4o[30]. This pattern suggests that stronger models already internalize some edge case reasoning from pretraining data, while weaker models depend more heavily on explicit prompt guidance. Research on class-level code generation has similarly found that model capability rankings can shift depending on task formulation [31].

B. Practical Cost-Benefit Considerations

The practical implications of the cross-model sensitivity analysis are noteworthy for resource-constrained practitioners. Qwen2.5-Coder-32B at L4 (EPR = 65.7%) approaches the performance of GPT-4o at L2 (EPR = 63.7%), suggesting that careful prompt engineering can partially substitute for model capability. Given that open-source model inference costs are typically an order of magnitude lower than closed-source API costs, investing in L4-level prompt preparation may be economically preferable to upgrading the underlying model. Benchmarks requiring complex multi-step reasoning have shown that even the strongest models leave substantial room for improvement [32], reinforcing the value of prompt-level interventions.

Table 5. Statistical Significance of Pairwise Specificity Level Comparisons on EPR

Comparison	Mean Δ EPR (pp)	Wilcoxon p-value	Cliff's Delta	Effect Size
L1 vs. L2	+5.4	< 0.01	0.37	Medium
L2 vs. L3	+5.8	< 0.01	0.41	Medium
L3 vs. L4	+4.7	< 0.05	0.28	Small
L1 vs. L4	+15.9	< 0.001	0.62	Large

Data source: Wilcoxon signed-rank tests with Bonferroni correction computed across all 164 problems, pooled over four LLMs. Effect sizes interpreted per standard thresholds: small ($|d| < 0.33$), medium ($0.33 \leq |d| < 0.47$), large ($|d| \geq 0.47$).

5. Discussion and Future Work

5.1. Practical Implications

The experimental results carry actionable implications for software development workflows that rely on LLM code generation. The finding that the L2-to-L3 transition (adding concrete examples) produces the largest single-step EPR gain suggests that practitioners should prioritize including representative input-output examples in their prompts. This intervention is both low-cost, as it requires only a few seconds of additional prompt preparation, and high-impact, as it triggers models to reason about the mapping between inputs and outputs rather than relying on pattern completion from the function signature alone. The additional gain from L3 to L4 (adding explicit edge case hints) is statistically significant yet smaller in magnitude, indicating diminishing returns from further specificity investment. A pragmatic strategy is to use L3-level prompts as the default and elevate to L4 only for safety-critical or edge-case-sensitive functions where the additional annotation effort is justified.

The cross-model analysis adds a resource-allocation dimension to these guidelines. Teams using open-source models with limited parameter counts stand to gain proportionally more from prompt specificity investment than teams already using the strongest available closed-source models. The finding that Qwen2.5-Coder-32B at L4 approaches GPT-4o at L2 suggests a concrete trade-off: organizations can achieve comparable edge case robustness at lower inference cost by combining a lighter model with higher prompt quality. The EPR metric introduced in this study offers a practical tool for teams that wish to audit the edge case robustness of their LLM-generated code. By running the EvalPlus augmented test suite and computing EPR alongside pass@1, teams can identify functions where generated code passes basic tests yet fails on boundary conditions, enabling targeted manual review.

5.2. Limitations and Future Directions

This study has several limitations that define opportunities for future investigation. The evaluation is restricted to Python and to function-level problems drawn from HumanEval+, which may not reflect the complexity of real-world software engineering tasks involving multi-file projects or library dependencies. Extending the analysis to class-level benchmarks and multilingual settings represents a natural next step. The prompt construction process, while guided by a codebook and validated by inter-annotator agreement, involves subjective judgment in determining what constitutes a meaningful edge case hint. Future work could explore automated prompt augmentation techniques that inject edge case specifications based on static analysis. The four-level specificity gradient represents a coarse discretization of a continuous space. Finer-grained manipulations would yield a more detailed understanding of the prompt-performance relationship. The study is also limited by the inherent non-determinism of LLM APIs, which, despite greedy decoding and majority voting, may introduce residual noise.

References

[1]. Chen, M., Tworek, J., Jun, H., Yuan, Q., Pinto, H. P. O., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., … Zaremba, W. (2021). Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374.

[2]. Liu, J., Xia, C. S., Wang, Y., & Zhang, L. (2023). Is your code generated by ChatGPT really correct? Rigorous evaluation of large language models for code generation. In Advances in Neural Information Processing Systems 36 (NeurIPS 2023).

[3]. Rubei, R., Nguyen, P. T., Bianchi, M., Pham, H. V., & Di Ruscio, D. (2025). Guidelines to prompt large language models for code generation: An empirical characterization. arXiv preprint arXiv:2601.13118.

[4]. Fakhoury, S., Chakraborty, S., Musuvathi, M., & Lahiri, S. K. (2024). LLM-based test-driven interactive code generation: User study and empirical evaluation. IEEE Transactions on Software Engineering, 50(9), 2424–2441.

[5]. Austin, J., Odena, A., Nye, M., Bosma, M., Michalewski, H., Dohan, D., Jiang, E., Cai, C., Terry, M., Le, Q., & Sutton, C. (2021). Program synthesis with large language models. arXiv preprint arXiv:2108.07732.

[6]. Cassano, F., Gouwar, J., Nguyen, D., Nguyen, S., Phipps-Costin, L., Pinckney, D., Yee, M.-H., Zi, Y., Anderson, C. J., Feldman, M. Q., Guha, A., Greenberg, M., & Jangda, A. (2023). MultiPL-E: A scalable and polyglot approach to benchmarking neural code generation. IEEE Transactions on Software Engineering, 49(7), 3456–3472.

[7]. Jain, N., Han, K., Gu, A., Li, W.-D., Yan, F., Zhang, T., Wang, S., Solar-Lezama, A., Sen, K., & Stoica, I. (2024). LiveCodeBench: Holistic and contamination free evaluation of large language models for code. arXiv preprint arXiv:2403.07974.

[8]. Riddell, M., Ni, A., & Cohan, A. (2024). Quantifying contamination in evaluating code generation capabilities of language models. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL 2024) (pp. 4268–4283).

[9]. Li, J., Li, G., Li, Y., & Jin, Z. (2025). Structured chain-of-thought prompting for code generation. ACM Transactions on Software Engineering and Methodology, 34(1), 1–28.

[10]. Le, H., Chen, H., Saha, A., Gokul, A., Sahoo, D., & Joty, S. (2024). CodeChain: Towards modular code generation through chain of self-revisions with representative sub-modules. In Proceedings of the 12th International Conference on Learning Representations (ICLR 2024).

[11]. Han, M., Jin, Z., & Zou, D. (2023). Comparative Evaluation of Self-Supervised Pretraining Strategies for Few-Shot Medical Image Analysis. Artificial Intelligence and Machine Learning Review, 7(2), 23-42

[12]. Wang, Z., Zhou, Z., Song, D., Huang, Y., Chen, S., Ma, L., & Zhang, T. (2021). Towards understanding the characteristics of code generation errors made by large language models. In Proceedings of the 47th IEEE/ACM International Conference on Software Engineering (ICSE 2022).

[13]. DeepSeek-AI, Zhu, Q., Guo, D., Shao, Z., Yang, D., Wang, P., Xu, R., Wu, Y., Li, Y., Gao, H., Bi, J., Xiao, Y., Lu, J., Huang, F., Zhang, H., Peng, R., & Luo, W. (2024). DeepSeek-Coder-V2: Breaking the barrier of closed-source models in code intelligence. arXiv preprint arXiv:2406.11931.

[14]. Ouyang, S., Zhang, J. M., Harman, M., & Wang, M. (2020). An empirical study of the non-determinism of ChatGPT in code generation. ACM Transactions on Software Engineering and Methodology, 34(2), 1–32.

[15]. Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. In Advances in Neural Information Processing Systems 36 (NeurIPS 2023).

[16]. Zhao, F., Zhang, M., Zhou, S., & Lou, Q. (2024). Application of deep reinforcement learning for cryptocurrency market trend forecasting and risk management.

[17]. Chen, Y., Chen, Z., & Zou, D. (2021). CarbonShift: Harnessing Grid Carbon Variability for Geo-Distributed Workload Scheduling. Artificial Intelligence and Machine Learning Review, 6(4), 18-31.

[18]. Li, Z., & Chen, Z. (2021). Performance Evaluation of Prompt Generation Strategies for AI Agents in Online Programming Education. Journal of Advanced Computing Systems, 5(9), 14-27.

[19]. Zhang, H., & Shi, W. (2024). Comparative Evaluation of Automated Detection Approaches for Identifying Implicit Compliance Violations in Cross-border Commercial Contract Clauses. Artificial Intelligence and Machine Learning Review, 7(2), 1-22.

[20]. Olausson, T. X., Inala, J. P., Wang, C., Gao, J., & Solar-Lezama, A. (2024). Is self-repair a silver bullet for code generation? In Proceedings of the 12th International Conference on Learning Representations (ICLR 2024).

[21]. Du, X., Liu, M., Wang, K., Wang, H., Liu, J., Chen, Y., Feng, J., Sha, C., Peng, X., & Lou, Y. (2024). Evaluating large language models in class-level code generation. In Proceedings of the 46th IEEE/ACM International Conference on Software Engineering (ICSE 2024).

[22]. Zhuo, T. Y., Vu, M. C., Chim, J., Hu, H., Yu, W., Widyasari, R., Yusuf, I. N., Zhan, H., He, J., Paul, I., Purber, S., Luo, B., Shi, T., Hust, W. G., Tran, N., Le, D., Duy, N., Hoang, N., Long, B., … Xia, C. S. (2020). BigCodeBench: Benchmarking code generation with diverse function calls and complex instructions. In Proceedings of the 13th International Conference on Learning Representations (ICLR 2020).

[23].]Zou, D., Chen, Z., & Ling, Z. (2019). A Comparative Evaluation of Deep Learning Paradigms for Low-Light Image Enhancement: From CNNs to Diffusion Models. Journal of Computing Innovations and Applications, 3(2), 85-95.