



Delft University of Technology

Document Version

Final published version

Licence

CC BY-NC-ND

Citation (APA)

Popescu, R. M. (2026). Reliable and Developer-Aligned Evaluation of Agents for Software Engineering. In S. H. Tan, & F. Khomh (Eds.), *FSE Companion 2026 - Proceedings of the 34th ACM International Conference on the Foundations of Software Engineering* (pp. 43-44). Association for Computing Machinery (ACM).
<https://doi.org/10.1145/3803437.3804877>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Reliable and Developer-Aligned Evaluation of Agents for Software Engineering

Razvan Mihai Popescu
r.m.popescu@tudelft.nl
Delft University of Technology
Delft, The Netherlands

Abstract

Large language models are rapidly moving towards closing the development cycle, transitioning from simple assistive companions to autonomous contributors deeply embedded into collaborative development environments. Despite their accelerated adoption, existing evaluation techniques are limited due to their fragmented nature and distorted projection of true model capabilities, often obtained from hypothetical syntactic scenarios. This research aims to bridge this gap by providing a comprehensive evaluation methodology for LLM-powered agents that is grounded in real-world software development practice. Our evaluation approach focuses on contamination-awareness, in-the-wild agentic behavior assessment, and trajectory-aware benchmarks and metrics capturing realistic coding contexts, human-aligned behavior, and model failure modes.

CCS Concepts

• **Software and its engineering** → **Software creation and management**.

Keywords

Large Language Models (LLMs), Autonomous AI Agents, Software Engineering, Evaluation, Benchmarking, Data Contamination

ACM Reference Format:

Razvan Mihai Popescu. 2026. Reliable and Developer-Aligned Evaluation of Agents for Software Engineering. In *34th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (FSE Companion '26)*, July 05–09, 2026, Montreal, QC, Canada. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3803437.3804877>

1 Introduction

Given their strong reasoning and abstraction capabilities, Large Language Models (LLMs) have been widely employed in the context of Software Engineering (SE), which has become a prominent application domain. These models have shown a remarkable capacity to understand the structured nature of source code and achieve strong performance in various code-related tasks, leading to an increasing reliance on empirical evaluations to assess their abilities and limitations.

However, current evaluation practices are often unreliable, difficult to reproduce, and poorly aligned with real-world developer

needs. Many works strongly rely on convenience-based metrics borrowed from NLP, such as BLEU, METEOR, or ROUGE, which fail to capture functional validity and developer utility [2, 7]. At the same time, many of these metrics are inconsistently used across different coding tasks with minimal to no adaptation, while their coding variants, such as CodeBLEU, or embedding-based versions, such as CodeBERTScore, have shown to perform on par with general translation metrics, while struggling to distinguish noise from genuine model predictions [4]. Furthermore, these metric-level limitations are amplified by flaws in widely used benchmarks. Various benchmarks suffer from issues such as saturation, data contamination, incorrect ground truths, and unrealistic context scenarios, which not only distort the reported performance, but also hinder the comparability and reproducibility of the studies [3]. Lastly, although LLMs are increasingly being used as evaluation judges, their assessments are prone to biases and hallucinations, often diverging from human judgements and further compromising the validity of evaluation protocols [8].

Nevertheless, these evaluation practices are still the norm in the community, despite often creating an illusion of model competence and resulting in misleading conclusions and potentially harmful downstream effects. The recent emergence of coding agents adds a new dimension to this evaluation landscape, as their autonomous and tool-mediated interactions with software systems challenge the assumptions of traditional one-shot evaluation setups. Therefore, there is a growing disparity between how these models are evaluated in research settings and how they are actually used by developers in practice.

This research aims to bridge this gap by developing a principled, multi-dimensional evaluation approach for LLM-powered agents in software engineering. Through systematic analysis of existing evaluation practices, empirical studies of agent behavior in real-world repositories, and the design of contamination-aware benchmarks targeting software evolution and maintainability, this work seeks to contribute towards a community-facing evaluation standard that enables reliable, reproducible, and developer-relevant assessment of LLM-based systems.

2 Related Work

Popular benchmarks such as HumanEval, MBPP, and Defects4J provide reproducible measures of functional correctness, yet are known to suffer from saturation, data contamination, and limited contextual realism issues (e.g., method-level operation) increasingly evident as frontier models, including GPT-o1 or Qwen-Coder¹, report near-ceiling performance on these tasks [3, 6]. To move beyond surface-level metrics, some studies employ LLMs as judges to score

¹<https://evalplus.github.io/leaderboard.html>



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

FSE Companion '26, Montreal, QC, Canada

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2636-1/2026/07

<https://doi.org/10.1145/3803437.3804877>

or rank generated code [8], while others rely on controlled user studies measuring productivity or developer satisfaction, complementing these benchmarks [1]. However, LLM-based judging raises concerns around bias and reproducibility, while human evaluations not only remain costly and difficult to scale but suffer from low external validity, failing to capture the complexity and pace of real-world development environments.

On the other hand, coding agents not only add a layer of autonomy and persistence over their precedents, but they are also capable of adjusting their behavior in response to observed effects while maintaining progress towards a user goal. Such agents are now able to perform multi-step task planning, take actions, run self-evaluations, and coordinate over different tools [5]. Current agentic benchmarks such as SWE-Bench, SWT-Bench, or AgentBench evaluate these agents under controlled conditions in different contexts, including issue resolution, bug reproduction, and decision making. However, these practices only capture narrow aspects of agent behavior, neglecting real-world collaboration and the ever-changing nature of development environments.

3 Research Plan

We outline three steps towards robust LLM evaluation in SE: a systematic review, in-the-wild assessment of agentic behavior, and benchmarking in real-world development workflows.

3.1 Evaluation Landscape

We first establish the empirical and methodological foundation for rethinking how LLMs should be evaluated in SE. We conducted a systematic literature review by synthesizing 279 peer-reviewed papers on 26 coding tasks in order to provide a comprehensive overview of the LLM4Code evaluation landscape, along with its limitations. We proposed a taxonomy of evaluation setups, audited the datasets and benchmarks currently in use, and tracked model adoption across the literature. By exposing the disconnect between performance scores, functional validity, and contextual clues, we establish a roadmap for rigorous, human-aligned, and functionally-sound evaluation frameworks. Our main research questions are as follows.

- RQ1:** Which LLMs have been evaluated in the context of code-related tasks?
- RQ2:** What datasets are driving the evaluation of LLMs across code-related tasks?
- RQ3:** What evaluation techniques have been applied to assess LLMs in code-related tasks

3.2 In-the-Wild Evaluation

The emergence of specialized coding agents powered by LLMs, including OpenAI Codex, GitHub Copilot, Claude Code, or Google Jules, opens up a new evaluation paradigm due to their tight operation within development environments and manipulation of version control systems. By exploiting their characteristic tell-tale signatures, we evaluate coding agents in open-source collaborative repositories, moving beyond the synthetic environments of conventional evaluation pipelines towards "in-the-wild" evaluations reflecting real coding workflows. We perform a thorough analysis

of both agentic and human development activities, along with a longitudinal investigation of their impact on codebase maintainability together with their failure points, and release a large-scale dataset of agentic and human contributions to support further research.

- RQ1:** What is the difference between agent-authored and human-authored activity in shaping collaboration and development progress?
- RQ2:** How do agent-authored contributions influence the trajectory of code maintenance over time compared to human-authored ones?

3.3 Change-Responsive Evaluation

In practice, software development typically unfolds through evolving requirements, iterative changes, and collaborative decision-making, shaping the trajectory of codebases over time. Based on these observations, this study introduces a multilingual, contamination-aware, benchmark for end-to-end issue resolution in evolving software repositories. Our main research questions are as follows.

- RQ1:** How do agents' performance patterns evolve when iteratively responding to feedback in complex development tasks?
- RQ2:** Are there systematic differences in behavior or failure modes across distinct parts of a software system under repeated modifications?

4 Conclusion

We propose a methodology for reliably evaluating LLM-powered agents in real-world development, focusing on human-aligned behavior, evolving environments, and failure resolution. We first systematize existing evaluation practices to identify limitations, then study in-the-wild agent behavior to uncover realistic challenges and failure modes, and finally use these insights to design benchmarks and metrics that better reflect developer intent and practice.

References

- [1] Joel Becker, Nate Rush, Elizabeth Barnes, and David Rein. 2025. Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity. arXiv:2507.09089 [cs.AI] <https://arxiv.org/abs/2507.09089>
- [2] Mikhail Evtikhiev, Egor Bogomolov, Yaroslav Sokolov, and Timofey Bryksin. 2023. Out of the BLEU: How should we assess quality of the Code Generation models? *J. Syst. Softw.* 203, C (Sept. 2023), 17 pages. doi:10.1016/j.jss.2023.111741
- [3] Xing Hu, Feifei Niu, Junkai Chen, Xin Zhou, Junwei Zhang, Junda He, Xin Xia, and David Lo. 2025. Assessing and Advancing Benchmarks for Evaluating Large Language Models in Software Engineering Tasks. arXiv:2505.08903 [cs.SE] <https://arxiv.org/abs/2505.08903>
- [4] Jonathan Katzy, Yongcheng Huang, Gopal-Raj Panchu, Maksym Ziemlewski, Paris Loizides, Sander Vermeulen, Arie van Deursen, and Maliheh Izadi. 2025. A Qualitative Investigation into LLM-Generated Multilingual Code Comments and Automatic Evaluation Metrics. arXiv:2505.15469 <https://arxiv.org/abs/2505.15469>
- [5] Junwei Liu, Kaixin Wang, Yixuan Chen, Xin Peng, Zhenpeng Chen, Lingming Zhang, and Yiling Lou. 2025. Large Language Model-Based Agents for Software Engineering: A Survey. arXiv:2409.02977 [cs.SE] <https://arxiv.org/abs/2409.02977>
- [6] Martin Riddell, Ansong Ni, and Arman Cohan. 2024. Quantifying Contamination in Evaluating Code Generation Capabilities of Language Models. arXiv:2403.04811 [cs.SE] <https://arxiv.org/abs/2403.04811>
- [7] Devjeet Roy, Sarah Fakhoury, and Venera Arnaoudova. 2021. Reassessing automatic evaluation metrics for code summarization tasks. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 1105–1116. doi:10.1145/3468264.3468588
- [8] Ruiqi Wang, Jiyu Guo, Cuiyun Gao, Guodong Fan, Chun Yong Chong, and Xin Xia. 2025. Can LLMs Replace Human Evaluators? An Empirical Study of LLM-as-a-Judge in Software Engineering. *Proceedings of the ACM on Software Engineering* 2, ISSTA (June 2025), 1955–1977. doi:10.1145/3728963