Fine-Tuning Impact on Qwen2.5 Code Generation Across HumanEval and MBPP Benchmarks

Assignee Research

May 31, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: How does the fine-tuning process for Qwen2.5 affect its performance on code generation benchmarks like HumanEval and MBPP compared to models trained on smaller pre-training datasets. We introduce self-invoking code generation, a new task designed to evaluate the progressive reasoning and problem-solving capabilities of LLMs. In this task, models are presented with a base problem and a related, more complex problem. 15 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: HumanEval Pro and MBPP Pro: Evaluating Large Language Models on Self-invoking Code Generation. Research question: How does the fine-tuning process for Qwen2.5 affect its performance on code generation benchmarks like HumanEval and MBPP compared to models trained on smaller pre-training datasets?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.0/10.

3 Results

4 papers retrieved. 15 claims extracted; 2 independently verified. Quality review score: 6.0/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
o1-mini achieves 96.2% pass@1 on HumanEval but only 76.2% on HumanEval Pro	✓	0.24
Instruction-tuned models are less efficient on self-invoking code generation than traditional code generation tasks	✓	0.33
HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021) serve as fundamental benchmarks, focusing on Python functio	×	0.04
Several benchmarks have expanded code evaluation benchmarks to encompass multiple programming languages (Zheng et al., 2	×	0.04
Benchmarks have expanded to include complex tasks like program repair (Haque et al., 2022; Jiang et al., 2023; Muennigho	×	0.04
Benchmarks have expanded to include dynamic problem sets (Jain et al., 2024)	×	0.04
Benchmarks have expanded to include code reasoning through code summarization (Barone and Sennrich, 2017; Hasan et al.,	×	0.05
Deepseek-V2.5 is used to generate self-invoking problems, candidate solutions, and test inputs	×	0.07
Generated solutions are executed with test inputs in a controlled Python environment to obtain ground truth outputs	×	0.00
An iterative method involving Python execution check and manual review ensures all test cases pass successfully	×	0.01
Final execution results are used to construct complete test cases with assert command	×	0.03
Qwen2.5-Coder-7B-base has a pass rate of 59.6% on HumanEval Pro and 38.6% on MBPP Pro	×	0.10
Qwen2.5-Coder-7B-instruct has a pass rate of 64.9% on HumanEval Pro and 35.1% on MBPP Pro	×	0.08
DeepseekCoder-33B-base has a pass rate of 71.9% on HumanEval Pro and 38.6% on MBPP Pro	×	0.11
DeepseekCoder-33B-instruct has a pass rate of 80.7% on HumanEval Pro and 43.9% on MBPP Pro	×	0.09

References

- http://arxiv.org/abs/1905.03197v3
- http://arxiv.org/abs/2001.11314v3
- http://arxiv.org/abs/2412.21199v2