

Mistral-Large-2 Self-Invoking Code Alignment and Cross-Domain Performance in MBPP Pro

Assignee Research

May 30, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does the alignment of Mistral-Large-2’s self-invoking code generation affect its performance on cross-domain tasks (e.g., math vs. string manipulation) in MBPP Pro, and can fine-tuning improve. 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. 12 claims were extracted from source literature; 3 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 5.2/10. 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 alignment of Mistral-Large-2’s self-invoking code generation affect its performance on cross-domain tasks (e.g., math vs. string manipulation) in MBPP Pro, and can fine-tuning improve generalization?.

2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 5.2/10.

3 Results

15 papers retrieved. 12 claims extracted; 3 independently verified. Quality review score: 5.2/10.

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.	✓	0.19
o1-mini achieves 76.2% pass@1 on HumanEval Pro.	✓	0.21
Instruction-tuned models are less efficient on self-invoking code generation than on traditional code generation tasks.	✓	0.26
HumanEval and MBPP serve as fundamental benchmarks focusing on Python function completion tasks with test-driven evaluation	×	0.08
Deepseek-V2.5 was used to generate self-invoking problems, candidate solutions, and test inputs for the benchmark construction	×	0.04
OpenCoder-8B-base achieved a score of 56.1 on the Base Problem and 10.5 on the Self-invoking Problem.	×	0.09
OpenCoder-8B-instruct achieved a score of 75.4 on the Base Problem and 22.8 on the Self-invoking Problem.	×	0.08
DeepseekCoder-6.7B-base achieved a score of 59.6 on the Base Problem and 35.1 on the Self-invoking Problem.	×	0.09
DeepseekCoder-6.7B-instruct achieved a score of 56.1 on the Base Problem and 35.1 on the Self-invoking Problem.	×	0.09
Qwen2.5Coder-7B-instruct achieved a score of 64.9 on the Base Problem and 35.1 on the Self-invoking Problem.	×	0.08
DeepseekCoder-33B-instruct achieved a score of 80.7 on the Base Problem and 43.9 on the Self-invoking Problem.	×	0.08
Yi-Coder-9B-Chat achieved a score of 66.7 on the Base Problem and 31.6 on the Self-invoking Problem.	×	0.07

References

- <http://arxiv.org/abs/2508.11281v3>
- <http://arxiv.org/abs/2306.08568v2>
- <http://arxiv.org/abs/2412.21199v2>