

Frontier Large Language Models in Mathematical Reasoning, Code Generation, and Scientific Knowledge

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: Comprehensive comparison of frontier large language models on mathematical reasoning code generation and scientific knowledge v17. 7 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: MR-GSM8K: A Meta-Reasoning Benchmark for Large Language Model Evaluation. Research question: Comprehensive comparison of frontier large language models on mathematical reasoning code generation and scientific knowledge v17.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.7/10.

3 Results

14 papers retrieved. 7 claims extracted; 0 independently verified. Quality review score: 3.7/10.

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
The MR-Score metric consists of three sub-metrics: Matthews Correlation Coefficient (MCC) for binary classification of s	×	0.02
The MCC score ranges from -1 to +1, where -1 indicates total disagreement between prediction and observation, 0 suggests	×	0.02
The models evaluated include Qwen-v1.5-1.8B, Llama3-70B, Deepseek-v2-236B, WizardMath-v1.1-7B, MAmmoTH-70B, DeepseekMath	×	0.06
Each model was evaluated under a zero-shot setting to assess their ability to follow instructions and their mathematical	×	0.08
The inference temperature was set to zero across all models to ensure reproducibility and minimize variance.	×	0.04
In the context of this paper, negative values are interpreted as no better than random guesses, and 0 is set as the cut-	×	0.04
The second metric is the accuracy of the first-error-step prediction: $ACC_{step} = N_{correct_first_error_step} / N_{incorrect_s}$	×	0.01

References

- <http://arxiv.org/abs/2509.25160v1>
- <http://arxiv.org/abs/2405.07551v1>
- <http://arxiv.org/abs/2312.17080v4>