

Current Language Model Benchmark Limitations in Reasoning Evaluation

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: What are the limitations of current language model evaluation benchmarks for measuring reasoning v13. 10 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.1/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: ARS: Adaptive Reasoning Suppression for Efficient Large Reasoning Language Models. Research question: What are the limitations of current language model evaluation benchmarks for measuring reasoning v13.

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 4.1/10.

3 Results

4 papers retrieved. 10 claims extracted; 0 independently verified. Quality review score: 4.1/10.

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 study evaluates Qwen2.5-Math-1.5B-Instruct, Qwen2.5-Math-7B-Instruct, and DeepSeek-R1-Distill-Qwen-7B models.	×	0.01
The ARS algorithm sets a maximum token limit of 1200 tokens per response.	×	0.02
Reasoning reflection behaviors are triggered by specific keywords including 'Wait', 'But', and 'Alternatively'.	×	0.03
ARS operates through three core components: Multi-checkpoint certainty estimation, Progressive threshold adaptation, and	×	0.14
In FAST mode, the ARS policy uses 2 drafts with 10 tokens per draft.	×	0.01
In MOD mode, the ARS policy uses a budget of 64 tokens.	×	0.01
In the default mode, the ARS policy uses a <code>sc_k</code> value of 3.	×	0.01
ARS consistently achieves superior length reduction while maintaining competitive accuracy across all tested model scale	×	0.09
The objective of the ARS framework is to minimize expected output length $E[T]$ subject to an acceptable accuracy degradat	×	0.02
ARS establishes multiple checkpoints at regular intervals during generation to estimate model certainty through tentativ	×	0.04

References

- <http://arxiv.org/abs/1610.00031v1>
- <http://arxiv.org/abs/2510.00071v2>

- <http://arxiv.org/abs/2407.04973v1>