Sparse Mixture-of-Experts vs. Dense Transformers in Mathematical Reasoning Benchmarks

Assignee Research

June 6, 2026

Abstract

This report synthesises findings from 12 peer-reviewed papers addressing the following research question: How do sparse mixture-of-experts models compare to dense transformers on mathematical reasoning v11. 16 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.5/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Measuring Mathematical Problem Solving With the MATH Dataset. Research question: How do sparse mixture-of-experts models compare to dense transformers on mathematical reasoning v11.

2 Methodology

Systematic literature search across multiple databases yielded 12 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.5/10.

3 Results

12 papers retrieved. 16 claims extracted; 0 independently verified. Quality review score: 3.5/10.

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
Accuracy remains low even for the best models on the MATH dataset.	×	0.14
Accuracy is increasing very slowly with model size on the MATH dataset.	×	0.09
Pretraining on AMPS enables a small 0.1B parameter model to perform similarly to a large fine-tuned 13B parameter model.	×	0.03
Having models generate their own step-by-step solutions before producing an answer actually degrades accuracy.	×	0.09
Providing partial ground truth step-by-step solutions can improve performance.	×	0.04
Providing models with step-by-step solutions at training time also increases accuracy.	×	0.06
GPT-2 models tokenize numbers so that one digit is processed at a time.	×	0.02
T5's tokenizer removes many LATEX symbols, making its performance not competitive.	×	0.02
Models pretrain on AMPS for one epoch using AdamW with a batch size of 128 and a weight decay of 0.05.	×	0.02
During pretraining, Khan Academy data is upsampled by a factor of 5 and Mathematica data is downsampled by a factor of 2	×	0.01
Models are trained with 8 A100 GPUs, each requiring less than a day.	×	0.02
For GPT-2, beam search has a beam size of 20 when only generating the final answer, and a beam size of 10 when generatin	×	0.03
GPT-2 0.1B model achieves an average accuracy of 5.4% on the MATH dataset.	×	0.10
GPT-2 1.5B model achieves an average accuracy of 8.8% on the MATH dataset.	×	0.10
GPT-3 13B model achieves an average accuracy of 5.8% on the MATH dataset.	×	0.10
GPT-3 175B model achieves an average accuracy of 3.0% on the MATH dataset.	×	0.10

References

- http://arxiv.org/abs/2603.11114v1
- http://arxiv.org/abs/2103.03874v2
- http://arxiv.org/abs/2402.14800v2