Scaling Laws of Chain-of-Thought Reasoning in Large Language Models

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What are the scaling laws for chain-of-thought reasoning in large language models v18. 20 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Intrinsic Stability Limits of Autoregressive Reasoning: Structural Consequences for Long-Horizon Execution. Research question: What are the scaling laws for chain-of-thought reasoning in large language models v18.

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 4.2/10.

3 Results

15 papers retrieved. 20 claims extracted; 1 independently verified. Quality review score: 4.2/10.

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
The sample complexity required for learning or planning can grow exponentially with decision depth.	×	0.03
Exponential growth in sample complexity is typically attributed to combinatorial explosion in state or policy spaces.	×	0.03
Temporal abstraction and the options framework reduce effective decision depth through macro-actions or hierarchical dec	×	0.03
Theorem A exhibits mathematically similar exponential behavior but differs fundamentally in interpretation.	×	0.01
Collapse can arise even in strictly linear settings without branching, due to cumulative noise within the autoregressive	×	0.03
Chain-of-Thought (CoT) prompting encourages models to generate intermediate reasoning steps, significantly improving per	×	0.14
Tree-of-Thought (ToT) introduces branching exploration and backtracking mechanisms to evaluate multiple reasoning candid	×	0.03
Most existing work implicitly treats reasoning as an autoregressive sequence that can be extended indefinitely.	×	0.04
The physical or informational constraints governing the stability of long reasoning trajectories remain largely unformal	×	0.04
Semantically coherent intermediate steps may lead to logical drift or inconsistency as reasoning depth increases.	×	0.05
The success of CoT or ToT implicitly relies on segmentation and reset mechanisms within the reasoning process.	×	0.08
Another line of research represents reasoning as graph structures, such as directed graphs or directed acyclic graphs (D	×	0.12
The autoregressive update is modeled as a stochastic dynamical process: Zt+1 = f(Zt) + εt.	×	0.04
In long-horizon sequences, perturbations εt typically accumulate rather than cancel out, leading to a monotonic erosion	×	0.09
Theorem A analyzes long-horizon autoregressive reasoning under a set of idealized structural assumptions.	✓	0.16
Theorem A is a structural stability theorem under contraction-like dynamics.	×	0.07
Single long edge: noise accumulation erodes directionality.	×	0.04
Directionality (decision advantage) decays with	×	0.08

References

- http://arxiv.org/abs/2503.09567v5
- http://arxiv.org/abs/2410.03595v1
- http://arxiv.org/abs/2602.06413v1