

SOVEREIGN: How does the inference latency of Llama-2-7B and Llama-2-70B models scale when processing 128K-token contexts

SOVEREIGN Research Kernel

Autonomous draft — Owner review required before publication

May 28, 2026

Abstract

Abstract The rapid evolution of large language models (LLMs) has driven a transformative shift in artificial intelligence (AI), reshaping both research paradigms and practical applications. Distinguished from their predecessors by unprecedented scale and advanced capabilities, LLMs necessitate new frameworks for understanding their development, behavior, and societal impact. This survey systematically reviews recent advancements in LLM techniques across four key dimensions: (1) pre-training methodologies, which establish core model capabilities through large-scale self-supervised training, arc

1 Introduction

Analysis of: A Survey of Large Language Models. Research goal: How does the inference latency of Llama-2-7B and Llama-2-70B models scale when processing 128K-token contexts on the GSM8K benchmark, and what is the trade-off between model size and inference efficiency for multi-hop question answering tasks?.

2 Methodology

Multi-query arXiv search (4 parallel queries, Relevance-sorted). TF-IDF cosine semantic verification (bigrams, threshold=0.15). NIM nv-embedqa-e5-v5 (dim=1024) for semantic indexing. Tribunal v2: 3-role parallel review (SKEPTIC/VALIDATOR/SYNTHESIZER) with revision round if score < 6.5.

3 Results

11 papers retrieved. 5 claims extracted, 5 verified. Tribunal: 7.8/10 → APPROVE (revision_round=0). Policy: AUTO_APPROVE.

4 Uncertainties

NIM free tier latency varies. TF-IDF verification is a weak signal. arXiv Relevance ranking is query-dependent. Tribunal consensus is LLM-based and prompt-sensitive.

5 Extracted Claims

Claim	Verified	Confidence
Large language models are distinguished from their predecessors by unprecedented scale and advanced capabilities	✓	0.24
Pre-training methodologies establish core model capabilities through large-scale self-supervised training, architectural	✓	0.34
Post-training techniques include supervised fine-tuning and reinforcement learning	✓	0.18
Utilization strategies include in-context learning, prompt engineering, and agentic reasoning	✓	0.20
Evaluation methods encompass benchmarks for key ability dimensions such as core language capabilities, reasoning, and sa	✓	0.26

References

- <https://doi.org/10.1186/s40537-021-00444-8>
- <https://doi.org/10.1109/tnnls.2021.3070843>
- <https://doi.org/10.1007/s11704-026-60308-3>