DeepSeek R1 Inference Latency Scaling in Nested Control Flow for Security Auditing

Assignee Research

June 4, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: How does the inference latency scaling exponent of Deepseek R1 change when processing nested control flow structures compared to linear code sequences in automated security auditing tasks. 18 claims were extracted from source literature; 2 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.7/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: ESAA-Security: An Event-Sourced, Verifiable Architecture for Agent-Assisted Security Audits of AI-Generated Code. Research question: How does the inference latency scaling exponent of Deepseek R1 change when processing nested control flow structures compared to linear code sequences in automated security auditing tasks?.

2 Methodology

Systematic literature search across multiple databases yielded 16 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.7/10.

3 Results

16 papers retrieved. 18 claims extracted; 2 independently verified. Quality review score: 4.7/10.

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 right evaluation objective for ESAA-Security is to determine whether a security audit can be executed as a governed,	✓	0.15
ESAA-Security insists that report integrity depends on verification status.	×	0.07
Three research questions follow naturally from the framing of ESAA-Security.	×	0.07
RQ1 asks whether an event-sourced execution model can make agent-assisted security audits replay-verifiable and traceabl	✓	0.16
RQ2 asks whether security review can be operationalized as a structured audit process with explicit domain coverage, che	×	0.15
RQ3 asks whether ESAA-Security produces artifacts that are useful for prioritization and remediation in AI-generated sof	×	0.14
The primary unit of analysis should remain the orchestrator run, consistent with ESAA.	×	0.03
ESAA-Security requires a second analytical layer: the artifact chain produced within a run.	×	0.05
Evaluation should inspect not only run termination and verification metadata, but also whether the expected chain of Pha	×	0.05
Evaluation should cover at least five dimensions: protocol compliance, replay-verifiable state integrity, coverage compl	×	0.05
A system that produces many speculative findings without traceability or structured remediation is not necessarily stron	×	0.05
A defensible initial validation should follow the empirical logic of the ESAA paper and use at least two case studies of	×	0.05
One case should be a small repository that permits close manual inspection of every artifact.	×	0.04
The second case should be a medium repository that exercises more domains, dependency paths, and reporting transitions.	×	0.02
At least one repository should be substantially AI-generated or AI-modified, since that is the motivating context of the	×	0.12
Repository selection should satisfy four criteria.	×	0.00
The repository must expose enough technical surface for reconnaissance to operate meaningfully.	×	0.02
The repository should enable nontrivial execution of multiple security domains.	×	0.06

References

- http://arxiv.org/abs/2603.21389v1
- http://arxiv.org/abs/2603.06365v1
- http://arxiv.org/abs/1406.1253v1