

Synthetic Dialogue Diversity Effects on LLM Safety Scores in AdvBench Evaluations

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: How does varying the diversity of synthetic dialogue data during alignment pretraining influence the safety evaluation scores of large language models on the AdvBench dataset. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Latent Jailbreak: A Benchmark for Evaluating Text Safety and Output Robustness of Large Language Models. Research question: How does varying the diversity of synthetic dialogue data during alignment pretraining influence the safety evaluation scores of large language models on the AdvBench dataset?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.2/10.

3 Results

14 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.2/10.

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.

References

- <http://arxiv.org/abs/2307.08487v3>
- <http://arxiv.org/abs/2403.00867v3>
- <http://arxiv.org/abs/2402.11651v2>