

DPO-Aligned vs. SFT-Only Models on HateXplain for Low-Resource Language Varieties

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does the performance of DPO-aligned models on low-resource language varieties (e.g., Spanish vs. Galician) compare to SFT-only models when evaluated on the HateXplain benchmark, measured by. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: AraDiCE: Benchmarks for Dialectal and Cultural Capabilities in LLMs. Research question: How does the performance of DPO-aligned models on low-resource language varieties (e.g., Spanish vs. Galician) compare to SFT-only models when evaluated on the HateXplain benchmark, measured by F1-score across dialects?.

2 Methodology

Systematic literature search across multiple databases yielded 13 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 6.3/10.

3 Results

13 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.3/10.

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/2409.11404v3>
- <http://arxiv.org/abs/2106.16020v1>
- <http://arxiv.org/abs/2509.09055v1>