

# Direct Preference Optimization vs. Reinforcement Learning for Multilingual Counter-Speech Generation

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: How does Direct Preference Optimization compare to Reinforcement Learning from Human Feedback in improving multilingual counter-speech generation quality on the XCSG benchmark. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.2/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Northeastern Uni at Multilingual Counterspeech Generation: Enhancing Counter Speech Generation with LLM Alignment through Direct Preference Optimization. Research question: How does Direct Preference Optimization compare to Reinforcement Learning from Human Feedback in improving multilingual counter-speech generation quality on the XCSG benchmark?.

## 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 6.2/10.

## 3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.2/10.

## 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/2412.15453v1>
- <http://arxiv.org/abs/2506.02018v2>
- <http://arxiv.org/abs/2407.14477v4>