

# High-Entropy Synthetic Caption Noise and Vision-Language Model Robustness in CLIP Architectures

Assignee Research

June 9, 2026

## Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: What is the impact of high-entropy synthetic caption noise on the robustness of vision-language model alignment across diverse domain shifts in CLIP-based architectures. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: Toward a Holistic Evaluation of Robustness in CLIP Models. Research question: What is the impact of high-entropy synthetic caption noise on the robustness of vision-language model alignment across diverse domain shifts in CLIP-based architectures?.

## 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 3.3/10.

## 3 Results

15 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 3.3/10.

## 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/2410.01534v2>
- <http://arxiv.org/abs/2507.22398v3>
- <http://arxiv.org/abs/2306.07713v3>