

Frozen Backbone vs Fine-Tuned Adapters in Cross-Domain Robot Vision Tasks

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 16 peer-reviewed papers addressing the following research question: How do frozen backbone vision-language models perform relative to fully fine-tuned adapters on cross-domain robot vision tasks measured by zero-shot classification accuracy. 0 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 6.0/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Toward More Reliable Artificial Intelligence: Reducing Hallucinations in Vision-Language Models. Research question: How do frozen backbone vision-language models perform relative to fully fine-tuned adapters on cross-domain robot vision tasks measured by zero-shot classification accuracy?.

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 6.0/10.

3 Results

16 papers retrieved. 0 claims extracted; 0 independently verified. Quality review score: 6.0/10.

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/2306.02080v3>
- <http://arxiv.org/abs/2306.11400v2>
- <http://arxiv.org/abs/2512.07564v1>