Latent Action Space Dimensionality and Sample Efficiency in Unsupervised Imitation Learning for Robotic Manipulation

Assignee Research

June 7, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does varying the dimensionality of latent action spaces in unsupervised imitation learning models affect sample efficiency on complex robotic manipulation benchmarks. 9 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.6/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: CLAM: Continuous Latent Action Models for Robot Learning from Unlabeled Demonstrations. Research question: How does varying the dimensionality of latent action spaces in unsupervised imitation learning models affect sample efficiency on complex robotic manipulation benchmarks?.

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 4.6/10.

3 Results

13 papers retrieved. 9 claims extracted; 1 independently verified. Quality review score: 4.6/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
CLAM outperforms all baselines and nearly matches the performance of BC with expert data in both state- and image-based	×	0.05
CLAM improves upon the best baseline VPT by more than 2$\times$ average normalized return on the DMControl (locomotion) tasks a	×	0.11
Transformer-CLAM achieves performance close to or even better than that of BC-Expert which uses the same amount of privi	×	0.09
All variants of CLAM outperform the best baseline VPT [11].	×	0.05
CLAM outperforms state-of-the-art methods in the problem setting where only play data is available as action-labeled dat	✓	0.18
CLAM enables scalable learning from easy-to-collect, cheap play data [21] avoiding the need for expensive task-specific	×	0.05
CLAM is evaluated on DMControl, MetaWorld, and CALVIN environments without modification.	×	0.03
CLAM is benchmarked on locomotion tasks (Hopper and HalfCheetah) and manipulation tasks (Assembly, Bin Picking, Peg Inse	×	0.03
For DMControl tasks, normalized return is reported following [22].	×	0.02

References

- http://arxiv.org/abs/2505.04999v1
- http://arxiv.org/abs/2107.08398v1
- http://arxiv.org/abs/2105.06411v2