

Discrete vs. Continuous Latent Actions in CLAM for Robot Imitation Learning

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 4 peer-reviewed papers addressing the following research question: How does the discrete latent action formulation in CLAM compare to continuous approaches in terms of sample efficiency and final success rates on the CALVIN benchmark when trained with varying levels. 8 claims were extracted from source literature; 7 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 8.0/10. 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 the discrete latent action formulation in CLAM compare to continuous approaches in terms of sample efficiency and final success rates on the CALVIN benchmark when trained with varying levels of demonstration diversity?.

2 Methodology

Systematic literature search across multiple databases yielded 4 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 8.0/10.

3 Results

4 papers retrieved. 8 claims extracted; 7 independently verified. Quality review score: 8.0/10.

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.

5 Extracted Claims

Claim	Verified	Confidence
Learning robot policies using imitation learning requires collecting large amounts of costly action-labeled expert demon	✓	0.36
Existing methods struggle when applied to complex robot tasks requiring fine-grained motions.	✓	0.27
CLAM incorporates continuous latent action labels instead of discrete representations.	✓	0.26
CLAM jointly trains an action decoder to ensure that the latent action space can be easily grounded to real actions with	✓	0.30
The labeled examples can be collected from non-optimal play data, enabling CLAM to learn performant policies without acc	✓	0.37
CLAM significantly outperforms prior state-of-the-art methods on continuous control benchmarks in DMControl (locomotion)	✓	0.30
CLAM achieves a 2-3x improvement in task success rate compared to the best baseline.	✓	0.20
CLAM is demonstrated on a real WidowX robot arm.	×	0.14

References

- <https://openalex.org/W7161353976>
- <https://doi.org/10.48550/arxiv.2505.04999>
- <https://doi.org/10.1007/s11704-024-40663-9>