

How does varying the latent space dimensionality in continuous action models affect sample efficiency and success

Assignee Research

June 10, 2026

Abstract

Learning robot policies using imitation learning requires collecting large amounts of costly action-labeled expert demonstrations, which fundamentally limits the scale of training data. A promising approach to address this bottleneck is to harness the abundance of unlabeled observations-e.g., from video demonstrations-to learn latent action labels in an unsupervised way. However, we find that existing methods struggle when applied to complex robot tasks requiring fine-grained motions. We design continuous latent action models (CLAM) which incorporate two key ingredients we find necessary for 1

1 Introduction

This paper examines: CLAM: Continuous Latent Action Models for Robot Learning from Unlabeled Demonstrations. Research question: How does varying the latent space dimensionality in continuous action models affect sample efficiency and success rates on fine-grained manipulation tasks compared to discrete token-based imitation learning methods?.

2 Methodology

Systematic literature search across multiple databases yielded 11 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 4.5/10.

3 Results

11 papers retrieved. 11 claims extracted; 1 independently verified. Quality review score: 4.5/10.

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	✓	0.17
The data regime of CLAM enables scalable learning from easy-to-collect, cheap play data, avoiding the need for expensive	×	0.07
CLAM uses a feedforward dimension of 2048, 4 attention heads, a dropout of 0.1, and GeLU as the feedforward activation i	×	0.03
The CALVIN Transformer CLAM model uses 6 encoder layers, 6 decoder layers, a feedforward dimension of 2048, 8 attention	×	0.05
The MetaWorld environment has a maximum episode step of 100, a state dimension of 39, an action dimension of 4, an image	×	0.03
The CALVIN environment has a maximum episode step of 200, a state dimension of 39, an action dimension of 7, an image sh	×	0.02
The evaluation environments in simulation include locomotion tasks from the DMControl benchmark (Hopper and HalfCheetah)	×	0.08

References

- <http://arxiv.org/abs/2208.01897v1>
- <http://arxiv.org/abs/2505.04999v1>
- <http://arxiv.org/abs/1709.10087v2>