

# How do continuous latent action models compare to discrete token baselines in terms of sample efficiency when

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 l

## 1 Introduction

This paper examines: CLAM: Continuous Latent Action Models for Robot Learning from Unlabeled Demonstrations. Research question: How do continuous latent action models compare to discrete token baselines in terms of sample efficiency when fine-tuned on language model benchmarks such as HELM or BIG-bench, measured by convergence speed and final accuracy?.

## 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 3.8/10.

## 3 Results

15 papers retrieved. 13 claims extracted; 1 independently verified. Quality review score: 3.8/10.

## 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.	×	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
CLAM scales with  Dunlabeled  while supervised IDMs only scale with  Dlabeled .	×	0.02
CLAM can leverage vast, unstructured observation data to learn latent actions in an unsuper-vised manner.	×	0.11
CLAM enables scalable learning from easy-to-collect, cheap play data avoiding the need for expensive task-specific data	×	0.06
The Transformer-CLAM model uses 6 encoder layers, 6 decoder layers, a feedforward dimension of 2048, 8 attention heads,	×	0.02
The CALVIN Transformer-CLAM model uses 6 encoder layers, 6 decoder layers, a feedforward dimension of 2048, 8 attention	×	0.02
The MetaWorld environment has a max episode steps of 100, state dimension of 39, action dimension of 4, image shape of [	×	0.04
The CALVIN environment has a max episode steps of 200, state dimension of 39, action dimension of 7, image shape of [84,	×	0.03
The evaluation environments in simulation include locomotion tasks from the DMControl benchmark (Hopper and HalfCheetah)	×	0.03

## References

- <http://arxiv.org/abs/2110.06500v2>
- <http://arxiv.org/abs/2505.04999v1>
- <http://arxiv.org/abs/2602.08239v1>