Latent Action Token Granularity and Sample Efficiency in Continuous Latent Action Models

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 7 peer-reviewed papers addressing the following research question: How does the granularity of latent action tokens in continuous latent action models affect sample efficiency on the CALVIN benchmark. 20 claims were extracted from source literature; 1 was independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 4.5/10. 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 granularity of latent action tokens in continuous latent action models affect sample efficiency on the CALVIN benchmark?.

2 Methodology

Systematic literature search across multiple databases yielded 7 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

7 papers retrieved. 20 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.	×	0.08
CLAM improves around 2-3$\times$ success rate on the MetaWorld (manipulation) tasks compared to the best baseline VPT.	×	0.13
Transformer-CLAM achieves performance close to or even better than that of BC-Expert which uses the same amount of privi	×	0.08
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 can leverage vast, unstructured observation data to learn latent actions in an unsupervised manner.	×	0.10
CLAM enables scalable learning from easy-to-collect, cheap play data, avoiding the need for expensive task-specific data	×	0.06
BC-AL using action-labeled data does not perform well due to imitating suboptimal demonstrations.	×	0.09
In the image domain, transfer from the pretrained IDM image encoder might cause improvements in performance.	×	0.04
For state-based inputs, the additional difficulty introduced by not training on ground-truth actions could regularize th	×	0.02
Latent action models scale with the size of unlabeled data (|Dunlabeled|) while supervised IDMs only scale with the size	×	0.13
The problem setup assumes |Dlabeled| |Dunlabeled|.	×	0.01
VPT learns a suboptimal IDM due to the limited amount of labeled data.	×	0.03
The choice of method for learning is dependent on the specific data regime.	×	0.03
CLAM is evaluated on DMControl, MetaWorld, and CALVIN environments.	×	0.05
All evaluated domains are continuous control environments with a fixed episode length and no termination conditions.	×	0.06
The evaluation tasks include locomotion tasks (Hopper and HalfCheetah) and manipulation tasks (Assembly, Bin Picking, Pe	×	0.05
CLAM is also evaluated in CALVIN with the	×	0.05

References

- http://arxiv.org/abs/2505.04999v1
- http://arxiv.org/abs/2010.14680v2
- http://arxiv.org/abs/2112.03227v4