Diffusion-Based Trajectory Editing vs. Static Reweighting in Cross-Domain Few-Shot Reasoning

Assignee Research

June 8, 2026

Abstract

This report synthesises findings from 13 peer-reviewed papers addressing the following research question: How does diffusion-based trajectory editing in xTED compare to static reweighting methods on FewTrans reasoning tasks in terms of cross-domain adaptation accuracy. 13 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.1/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: xTED: Cross-Domain Adaptation via Diffusion-Based Trajectory Editing. Research question: How does diffusion-based trajectory editing in xTED compare to static reweighting methods on FewTrans reasoning tasks in terms of cross-domain adaptation accuracy?.

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 3.1/10.

3 Results

13 papers retrieved. 13 claims extracted; 0 independently verified. Quality review score: 3.1/10.

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
The WidowX robot and Airbot were used to collect 100 trajectories each for real-robot experiments.	×	0.02
Three manipulation tasks were built: (1) Pick up a red cup and place it on a silver pan (Cup); (2) Pick up a duck and pl	×	0.01
The task environments are highly stochastic, featuring randomly initialized object positions and poses, as well as distr	×	0.02
Trajectories were edited with latent observation encoded with DecisionNCE and reward-relevant modules were removed in th	×	0.04
Simulation experiments were conducted using the MuJoCo physics simulator.	×	0.01
Two source domains were constructed on Walker2d-v2 (WK) and HalfCheetah-v2 (HC) with introduced dynamics and morphology	×	0.04
Gravity was applied with 2$\times$ gravitational acceleration, friction coefficients were set to 0.25/0.5$\times$, and thigh size was	×	0.01
20k transitions were randomly selected from the Medium (M), Medium Replay (MR), and Medium Expert (ME) datasets in the D	×	0.02
The objective is to minimize the difference of the forward posterior and reverse conditional distribution across all den	×	0.02
The variance of the approximate reverse distribution is constructed as Σ = Σ = 2I.	×	0.02
A simplified surrogate loss is adopted instead of the distribution mean matching.	×	0.03
Diffusion-based image editing methods aim to modify visual styles while preserving semantic content in source images.	×	0.10
Diffusion editing methods achieve partial modifications by initiating the reverse process from an intermediate noise ste	×	0.04

References

- http://arxiv.org/abs/2409.08687v4
- http://arxiv.org/abs/2410.06883v5
- http://arxiv.org/abs/2003.02366v2