

xTED: Cross-Domain Adaptation via Diffusion-Based Trajectory Editing

Extended Abstract

Haoyi Niu^{*}
Tsinghua University
Beijing, China

Qimao Chen^{*}
Tsinghua University
Beijing, China

Tenglong Liu
National University of
Defense Technology
Changsha, China

Jianxiong Li
Tsinghua University
Beijing, China

Guyue Zhou
Tsinghua University
Beijing, China

Yi Zhang[✉]
Tsinghua University
Beijing, China

Jianming Hu[✉]
Tsinghua University
Beijing, China

Xianyuan Zhan[✉]
Tsinghua University
Beijing, China

ABSTRACT

Reusing cross-domain data is vital for decision-making tasks with limited target data, yet existing methods often rely on complex, inflexible model-level designs. To address this, we propose the Cross-Domain Trajectory EDiting (xTED) framework, which bridges domain gaps directly at the data level. Our proposed model architecture effectively captures the intricate dependencies among states, actions, and rewards, as well as the dynamics patterns within target data. Edited by adding noises and denoising with the pre-trained target diffusion model, source domain trajectories can be transformed to align with target domain properties while preserving original task semantic information. This process effectively corrects underlying domain gaps, enhancing state realism and dynamics reliability in source data, and allowing flexible integration with various single-domain and cross-domain downstream policy learning methods. Despite its simplicity, xTED demonstrates superior performance in extensive simulation and real-robot experiments. The full version can be found at <https://arxiv.org/abs/2409.08687>.

KEYWORDS

Cross-Domain, Diffusion-Based Editing, Robot Learning

ACM Reference Format:

Haoyi Niu^{*}, Qimao Chen^{*}, Tenglong Liu, Jianxiong Li, Guyue Zhou, Yi Zhang[✉], Jianming Hu[✉], and Xianyuan Zhan[✉]. 2026. xTED: Cross-Domain Adaptation via Diffusion-Based Trajectory Editing: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/BENP4894>

1 INTRODUCTION

Solving real-world tasks with reinforcement learning (RL) or imitation learning (IL) faces serious data scarcity issues [1, 13]. While incorporating data from source domains (e.g., simulation) is a promising solution [7–9], inevitable domain gaps—such as dynamics and

morphology discrepancies—often lead to negative transfer when such data is used directly [6].

Existing cross-domain methods typically rely on complex, domain-specific model designs, such as discriminators or specialized encoders [11, 12, 14]. These approaches often lack flexibility and hinder efficient data reuse across multiple domains. This limitation motivates us to rethink: *instead of adapting the policy learning process, can we directly bridge these gaps at the data level?*

Inspired by diffusion-based image editing [3, 5], which transforms visual styles while preserving content, we propose the **Cross-Domain Trajectory EDiting (xTED)** framework. However, applying this philosophy to decision-making is challenging because trajectory data consists of heterogeneous elements (states, actions, rewards) with intricate internal dependencies, unlike homogeneous pixels in images. To address this, xTED employs a novel diffusion architecture tailored for decision-making data. It features separate encoding/decoding for each element type to preserve physical meanings and introduces dependency structure modeling to capture dynamics relationships.

The xTED pipeline involves three simple steps: (1) train a diffusion model on target data to capture the target prior; (2) perturb source trajectories with noise and denoise them using the pre-trained model to align with target properties [2]; and (3) incorporate the edited data into downstream policy learning. By correcting domain gaps at the data level, xTED allows for flexible integration with various algorithms. Extensive experiments demonstrate that xTED consistently yields performance gains over baselines, particularly in challenging real-robot manipulation tasks.


2 MODEL ARCHITECTURE

To address the heterogeneity of decision-making data [10], xTED encodes states (s), actions (a), and rewards (r) separately using independent encoders and self-attention blocks. Critically, to capture MDP dynamics while respecting causal dependencies (where rewards stem from state-action pairs), we introduce a structured cross-attention mechanism:

$$h_k^{ts}, h_k^{ta} = \text{X-Attn}(h_k^s, h_k^a, \mathbf{e}_k^c), \text{X-Attn}(h_k^a, h_k^s, \mathbf{e}_k^c) \quad (1)$$

$$h_k^{tr} = \text{X-Attn}(h_k^r, [h_k^s, h_k^a], \mathbf{e}_k^c) \quad (2)$$

^{*}Equal Contribution: t6.da.thu@gmail.com, cqm24@mails.tsinghua.edu.cn.

[✉]Corresponding authors: {zhyi, hujm, zhanxianyuan@tsinghua.edu.cn}.
 This work is licensed under a Creative Commons Attribution International 4.0 License.

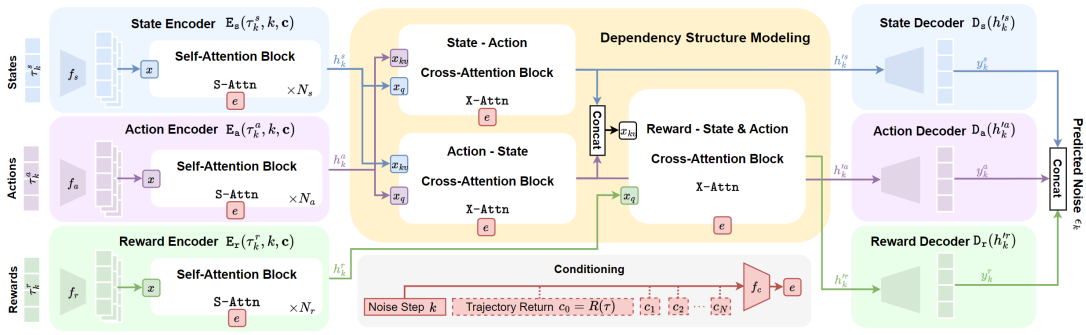


Figure 1: The model architecture.

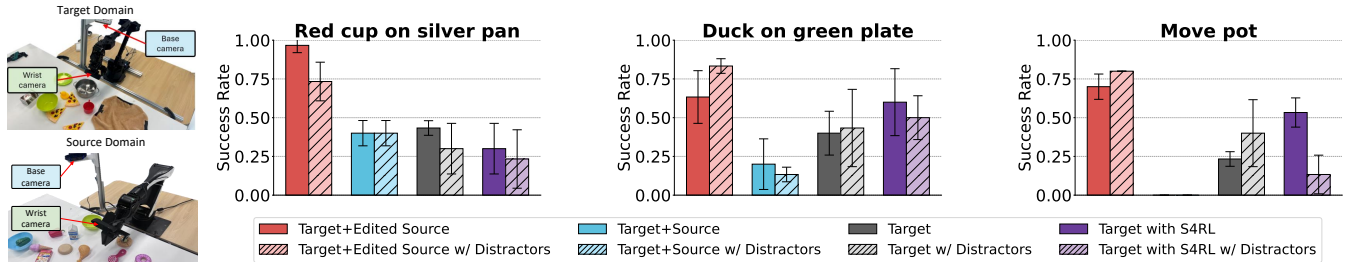


Figure 2: Real robot results and configuration.

This design, as shown in Fig. 1, models the mutual interaction between states and actions, while forcing rewards to query the joint state-action representation, effectively preventing unreasonable reverse dependencies.

Cross-Domain Editing Pipeline. The editing strategy follows a noise-then-denoise paradigm [5]. First, we train the diffusion model on **target data** to capture its trajectory distribution prior. Second, we obtain noised source trajectories $\tilde{\tau}_k$ from $\mathcal{N}(\tilde{\tau}_k; \tilde{\tau}_0, \sigma_k^2 \mathbf{I})$, which retains some information from the original source trajectories $\tilde{\tau}_0$. This procedure is governed by a ratio parameter $\kappa = \frac{k}{K} \in [0, 1]$, determining the level of noise addition. Lastly, we denoise $\tilde{\tau}_k$ using the learned target diffusion prior, aligning the edited trajectories $\hat{\tau}_0$ with the target domain. This process projects source skills onto the target manifold, blurring domain-specific biases while preserving semantic task information for downstream policy learning.

3 EXPERIMENTS

Real-robot. As shown in Fig. 2, using BC, **Target+Edited Source** overwhelmingly outperforms baselines across **Cup**, **Duck**, and **Pot** tasks. While adding raw source data causes severe negative transfer (e.g., success rate drops to 0% in *Pot*) and simple augmentation (S4RL) proves ineffective, xTED successfully leverages biased data, boosting success rates from 43% to **97%** in the *Cup* task.

Simulation. We integrate xTED with Implicit Q-Learning (IQL) [4]. As shown in Tab. 1, across 18 distinct task settings, xTED achieves the best or on-par performance in **all 18 tasks**. While direct augmentation with raw source data hurts performance in 5/18 tasks due to dynamics mismatches, xTED consistently provides meaningful augmentation, yielding an average improvement of **16.4%** over the Target-only baseline. Notably, in Medium-Replay datasets, xTED boosts scores by over 50% (Walker2d) and 20% (HalfCheetah).

Table 1: Average normalized scores for MuJoCo tasks on 20k target/source data (5 random seeds).

Target Data	Domain Gap	Src	Tgt	Tgt+Src	Δ	Tgt+Src (Edited)	Δ	
HalfCheetah	Med	Gravity	8.4±6.3	39.5±2.4	41.1±1.3	+4.1%	40.6±2.0	+2.8%
		Friction	14.4±3.6	39.5±2.4	27.9±7.2	-29.4%	41.2±1.6	+4.3%
		Thigh Size	-0.1±1.0	39.5±2.4	40.3±2.2	+2.0%	40.7±2.4	+3.0%
	Med-R	Gravity	8.4±6.3	26.2±3.5	29.5±2.9	+12.6%	31.3±2.9	+19.5%
		Friction	14.4±3.6	26.2±3.5	23.2±4.0	-11.5%	31.8±3.1	+21.4%
		Thigh Size	-0.1±1.0	26.2±3.5	29.0±3.7	+10.7%	33.0±3.0	+26.0%
	Med-E	Gravity	8.4±6.3	40.1±2.9	40.4±3.4	+0.7%	43.8±3.6	+9.2%
		Friction	14.4±3.6	40.1±2.9	27.1±5.4	-32.4%	43.2±3.0	+7.7%
		Thigh Size	-0.1±1.0	40.1±2.9	37.4±5.5	-6.7%	43.0±3.0	+7.2%
Walker2d	Med	Gravity	16.6±6.2	45.3±15.9	60.4±9.5	+33.3%	58.2±11.7	+28.5%
		Friction	8.2±1.4	45.3±15.9	47.0±12.8	+3.8%	54.5±13.7	+20.3%
		Thigh Size	7.6±3.5	45.3±15.9	49.6±12.2	+9.5%	58.9±11.7	+30.0%
	Med-R	Gravity	16.6±6.2	16.6±5.9	19.5±10.7	+17.5%	23.3±9.1	+40.3%
		Friction	8.2±1.4	16.6±5.9	17.4±6.2	+4.8%	25.9±9.1	+56.0%
		Thigh Size	7.6±3.5	16.6±5.9	18.0±6.7	+8.4%	25.9±9.1	+56.0%
	Med-E	Gravity	16.6±6.2	71.0±21.0	67.4±11.1	-16.9%	82.9±18.1	+16.8%
		Friction	8.2±1.4	71.0±21.0	75.1±18.3	+5.8%	74.0±24.5	+4.2%
		Thigh Size	7.6±3.5	71.0±21.0	77.0±19.3	+8.5%	81.0±21.4	+14.1%
Total		165.3	716.1	727.3	+1.6%	833.2	+16.4%	

4 CONCLUSION

We present xTED, a generic framework that reframes cross-domain policy adaptation as a data pre-processing task. By decoupling adaptation from policy learning, xTED is task- and domain-agnostic, compatible with various observation encoders and algorithms, and orthogonal to other adaptation methods. At its core, xTED introduces a specialized diffusion architecture that effectively models the complex dependencies within heterogeneous trajectory data. We show that incorporating source data edited with xTED consistently yields performance improvements over training solely on target data while directly adding unprocessed source data often results in significant performance degradation, particularly in real-robot manipulation tasks.

ACKNOWLEDGMENTS

This work is supported by Wuxi Research Institute of Applied Technologies, Tsinghua University, under Grant 20242001120, and Xiongan AI Institute.

REFERENCES

- [1] Jérémie Guiochet, Mathilde Machin, and Hélène Waeselynck. 2017. Safety-critical advanced robots: A survey. *Robotics and Autonomous Systems* 94 (2017), 43–52.
- [2] Yi Huang, Jiancheng Huang, Yifan Liu, Mingfu Yan, Jiayi Lv, Jianzhuang Liu, Wei Xiong, He Zhang, Shifeng Chen, and Liangliang Cao. 2024. Diffusion model-based image editing: A survey. *arXiv preprint arXiv:2402.17525* (2024).
- [3] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6007–6017.
- [4] Ilya Kostrikov, Ashvin Nair, and Sergey Levine. 2022. Offline Reinforcement Learning with Implicit Q-Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=68n2s9ZJWF8>
- [5] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *International Conference on Learning Representations*. https://openreview.net/forum?id=aBsCjcPu_tE
- [6] Haoyi Niu, Jianming Hu, Guyue Zhou, and Xianyuan Zhan. 2024. A Comprehensive Survey of Cross-Domain Policy Transfer for Embodied Agents. In *International Joint Conference on Artificial Intelligence*.
- [7] Haoyi Niu, Tianying Ji, Bingqi Liu, Haocheng Zhao, Xiangyu Zhu, Jianying Zheng, Pengfei Huang, Guyue Zhou, Jianming Hu, and Xianyuan Zhan. 2023. H2O+: An Improved Framework for Hybrid Offline-and-Online RL with Dynamics Gaps. *arXiv preprint arXiv:2309.12716* (2023).
- [8] Haoyi Niu, Shubham Sharma, Yiwen Qiu, Ming Li, Guyue Zhou, Jianming HU, and Xianyuan Zhan. 2022. When to Trust Your Simulator: Dynamics-Aware Hybrid Offline-and-Online Reinforcement Learning. In *Advances in Neural Information Processing Systems*. <https://openreview.net/forum?id=zXE8iFOZKw>
- [9] Open X-Embodiment et al. 2023. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. In *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition @ CoRL2023*. <https://openreview.net/forum?id=zraBtFgxT0>
- [10] William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 4195–4205.
- [11] Bradly C Stadie, Pieter Abbeel, and Ilya Sutskever. 2016. Third Person Imitation Learning. In *International Conference on Learning Representations*.
- [12] Guan Wang, Haoyi Niu, Desheng Zhu, Jianming Hu, Xianyuan Zhan, and Guyue Zhou. 2022. A Versatile and Efficient Reinforcement Learning Approach for Autonomous Driving. *NeurIPS Machine Learning for Autonomous Driving Workshop* (2022). https://github.com/imoneoi/carla_env
- [13] Xianyuan Zhan, Haoran Xu, Yue Zhang, Xiangyu Zhu, Honglei Yin, and Yu Zheng. 2022. DeepThermal: Combustion Optimization for Thermal Power Generating Units Using Offline Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [14] Qiang Zhang, Tete Xiao, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. 2020. Learning Cross-Domain Correspondence for Control with Dynamics Cycle-Consistency. In *International Conference on Learning Representations*.