

Novelty Assessment Report

Paper: Geometry-aware Policy Imitation

PDF URL: <https://openreview.net/pdf?id=ggofj6tyr3>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-04

Abstract

We propose a Geometry-Aware Policy Imitation (GPI) approach that rethinks imitation learning by treating demonstrations as geometric curves rather than collections of state-action samples. From these curves, GPI derives distance fields that give rise to two complementary control primitives: a progression flow that advances along expert trajectories and an attraction flow that corrects deviations. Their combination defines a controllable, non-parametric vector field that directly guides robot behavior. This formulation decouples metric learning from policy synthesis, enabling modular adaptation across low-dimensional robot states and high-dimensional perceptual inputs. GPI naturally supports multimodality by preserving distinct demonstrations as separate models and allows efficient composition of new demonstrations through simple additions to the distance field. We evaluate GPI in simulation and on real robots across diverse tasks. Experiments show that GPI achieves higher success rates than diffusion-based policies while running 20× faster, requiring less memory, and remaining robust to perturbations. These results establish GPI as an efficient, interpretable, and scalable alternative to generative approaches for robotic imitation learning.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Learning from Demonstrations via Geometric Distance Fields**

A total of **31 papers** were analyzed and organized into a taxonomy with **17 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Geometric Representation and Distance Field Methods**
- **Trajectory Representation and Encoding**
- **Optimal Transport and Divergence Minimization**
- **Reinforcement Learning-Based Imitation**
- **Visual and Perceptual Imitation Learning**
- **Domain-Specific Imitation Applications**
- **Non-Technical Studies**

Complete Taxonomy Tree

- Learning from Demonstrations via Geometric Distance Fields Survey Taxonomy
- Geometric Representation and Distance Field Methods
 - Distance Field-Based Policy Synthesis ★ (3 papers)
 - [0] Geometry-aware Policy Imitation (Anon et al., 2026) [View paper](#)
 - [4] Diff-ldf: Contact-aware model-based learning from visual demonstration for robotic manipulation via differentiable physics-based simulation and rendering (X Zhu, 2023) [View paper](#)
 - [14] DFields: Dynamic 3D Descriptor Fields for Zero-Shot Generalizable Robotic Manipulation (Y Wang, 2023) [View paper](#)
 - Geometric Constraint Inference (3 papers)
 - [8] Learning Rhythmic Trajectories With Geometric Constraints for Laser-Based Skincare Procedures (Anqing Duan, 2023) [View paper](#)
 - [23] Tracing curves in the plane: Geometric-invariant learning from human demonstrations. (Sri Harsha Turlapati, 2024) [View paper](#)
 - [29] Inferring geometric constraints in human demonstrations (Guru Subramani, 2022) [View paper](#)
 - Geometric Feature-Based Skill Learning (2 papers)
 - [6] Peg-in-hole assembly skill imitation learning method based on ProMPs under task geometric representation (Yajing Zang, 2023) [View paper](#)
 - [30] A Geometric Perspective on Visual Imitation Learning (Jun Jin, 2022) [View paper](#)
- Trajectory Representation and Encoding
 - Riemannian Manifold Trajectory Learning (2 papers)
 - [9] A structured prediction approach for robot imitation learning (Duan, 2024) [View paper](#)
 - [15] An approach for imitation learning on Riemannian manifolds (Martijn J.A. Zeestraten, 2017) [View paper](#)
 - Movement Primitive and Probabilistic Trajectory Methods (2 papers)
 - [3] A Trajectory Optimisation-Based Incremental Learning Strategy for Learning from Demonstration (Yuqi Wang, 2024) [View paper](#)
 - [17] Complex Robot Motion Generation via Context-Driven Sequencing of Movement Primitives (Edgar, 2025) [View paper](#)
- Optimal Transport and Divergence Minimization
 - Sinkhorn and Wasserstein Distance Methods (2 papers)
 - [11] Imitation Learning with Sinkhorn Distances (Georgios Papagiannis, 2022) [View paper](#)
 - [12] Adversarial Imitation Learning Based on Weighted Wasserstein Distance (Zhengzuo Qin, 2025) [View paper](#)

- Optimal Transport for Offline Imitation (1 papers)
- [25] Align Your Intent: Offline Imitation Learning via Optimal Transport (Buzun, 2024) [View paper](#)
- Reinforcement Learning-Based Imitation
 - Goal Proximity and Reward Shaping (2 papers)
 - [10] Generalizable imitation learning from observation via inferring goal proximity (Youngwoon Lee, 2021) [View paper](#)
 - [26] Expert Proximity as Surrogate Rewards for Single Demonstration Imitation Learning (Lan, 2024) [View paper](#)
 - Skill Diversity and Trajectory Matching (2 papers)
 - [13] Sd-pdmd: Deep reinforcement learning for robotic trajectory imitation (Yue Li, 2022) [View paper](#)
 - [19] Demonstration-Based Proximal Policy Optimization with Action Guidance (Zixuan Liu, 2022) [View paper](#)
- Visual and Perceptual Imitation Learning
 - Video-Based Demonstration Parsing (1 papers)
 - [16] One-shot learning for rapid generation of structured robotic manipulation tasks from 3D video demonstrations (Jaime Duque Domingo, 2025) [View paper](#)
 - Visuomotor Policy Learning (1 papers)
 - [2] Universal planning networks: Learning generalizable representations for visuomotor control (Aravind Srinivas, 2018) [View paper](#)
 - Trajectory Following in Visual Environments (1 papers)
 - [7] Adapting a world model for trajectory following in a 3d game (Ishida Shu, 2025) [View paper](#)
- Domain-Specific Imitation Applications
 - Autonomous Driving and Navigation (2 papers)
 - [5] Flow matching-based autonomous driving planning with advanced interactive behavior modeling (Tan Tianyi, 2025) [View paper](#)
 - [20] An imitation from observation approach for dozing distance learning in autonomous bulldozer operation (Ke You, 2022) [View paper](#)
 - Robotic Manipulation (1 papers)
 - [31] Methods for Teaching Diverse Robot Skills: Leveraging Priors, Geometry, and Dynamics (Asif, 2020) [View paper](#)
 - Multi-Agent and Collective Behavior (2 papers)
 - [18] Learning aggressive animal locomotion skills for quadrupedal robots solely from monocular videos (Liu Zhao, 2025) [View paper](#)
 - [24] Learning to Imitate Spatial Organization in Multi-robot Systems (Ramchurn, 2024) [View paper](#)
 - Spatio-Temporal Event Modeling (1 papers)
 - [1] Imitation learning of neural spatio-temporal point processes (Zhu, 2021) [View paper](#)
- Non-Technical Studies (4 papers)
 - [21] Distance to frontier: From imitation to innovation (Ayaz Zeynalov, 2025) [View paper](#)
 - [22] Neighborly Gains? The Green Spillover Effects of Lowâ€¢Carbon City Pilot Policy (H Wei, 2025) [View paper](#)
 - [27] Problems in the Difference-in-Distance measure of phonetic imitation (Bethany MacLeod, 2021) [View paper](#)
 - [28] Cognitive Control Structures in the Imitation Learning of Spatial Sequences and Rhythms-An fMRI Study. (Katrin Sakreida, 2019) [View paper](#)

Narrative

Core task: learning from demonstrations via geometric distance fields. The field organizes around several complementary perspectives on how to extract and reproduce demonstrated behaviors. Geometric Representation and Distance Field Methods emphasize spatial structure and collision-free motion synthesis, often leveraging signed distance functions or manifold-based encodings to capture task constraints. Trajectory Representation and Encoding focuses on compressing and generalizing demonstrated paths through probabilistic models, dynamic movement primitives, or neural embeddings. Optimal Transport and Divergence Minimization treats imitation as a distribution-matching problem, using Wasserstein distances or Sinkhorn divergences to align learned policies with expert data. Reinforcement Learning-Based Imitation blends demonstration data with trial-and-error optimization, while Visual and Perceptual Imitation Learning tackles the challenge of learning from raw sensory inputs. Domain-Specific Imitation Applications explore targeted use cases such as robotic manipulation, autonomous driving, or multi-agent coordination, each imposing unique geometric or temporal constraints.

Within the geometric branch, a small cluster of works explores how distance fields can directly guide policy synthesis. DFields[14] constructs neural distance representations to encode spatial constraints from demonstrations, while Diff-lfd[4] uses diffusion models to generate trajectories that respect learned geometric structure. Geometry-aware Policy Imitation[0] sits naturally alongside these methods, emphasizing the integration of geometric priors into the imitation pipeline to ensure physically plausible and collision-aware behavior. Compared to Trajectory Optimisation Incremental[3], which refines paths through iterative optimization, Geometry-aware Policy Imitation[0] more directly encodes spatial relationships into the policy architecture. This line of work addresses a central trade-off: balancing the expressiveness of learned representations with the interpretability and safety guarantees that explicit geometric reasoning provides, a question that remains active as methods scale to more complex environments and richer sensory modalities.

Related Works in Same Category

No comparison data available.

Contributions Analysis

Overall novelty summary. The paper proposes treating demonstrations as geometric curves and deriving distance fields to generate progression and attraction flows for robot control. It resides in the 'Distance Field-Based Policy Synthesis' leaf, which contains only three papers total, including this work and two siblings (DFields and Diff-lfd). This represents a relatively sparse research direction within the broader taxonomy of 31 papers across the field, suggesting the geometric distance field approach to policy synthesis remains an emerging area rather than a crowded subfield.

The taxonomy reveals that neighboring research directions include 'Geometric Constraint Inference' (extracting kinematic constraints from demonstrations) and 'Movement Primitive and Probabilistic Trajectory Methods' (encoding trajectories via dynamic movement primitives). The paper's approach diverges from these by directly synthesizing control policies from distance fields rather than extracting constraints or encoding trajectories probabilistically. It also differs from the 'Optimal Transport and Divergence Minimization' branch, which treats imitation as distribution matching rather than geometric curve following, and from 'Visual and Perceptual Imitation Learning', which focuses on high-dimensional sensory inputs rather than geometric structure.

The analysis examined zero candidate papers for all three contributions, meaning no literature search was conducted to identify potentially overlapping prior work. Without examining any candidates, the contribution-level statistics provide no evidence about whether the geometric distance field formulation, the decoupling of metric learning from policy synthesis, or the efficiency claims have substantial precedent. The absence of a literature search leaves the novelty assessment entirely dependent on the taxonomy structure and the two sibling papers in the same leaf, which address related but distinct aspects of geometric policy learning.

Given the limited search scope (zero candidates examined), this assessment reflects only the paper's position within a sparse taxonomy leaf and its relationship to two sibling works. The geometric distance field approach appears to occupy a relatively unexplored niche, but a comprehensive novelty evaluation would require examining a broader set of candidates from related leaves and potentially from outside the provided taxonomy structure.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Geometry-aware Policy Imitation (GPI) approach

Description: GPI represents expert demonstrations as geometric curves that induce distance fields in state space. These fields give rise to two complementary control primitives: a progression flow advancing along trajectories and an attraction flow correcting deviations. Their combination defines a controllable, non-parametric vector field that directly guides robot behavior.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 2: Modular formulation decoupling metric learning from policy synthesis

Description: The approach separates metric learning (defining how states are represented and compared) from behavior synthesis (constructing policies from distance and flow fields). This decoupling enables flexible adaptation across low-dimensional robot states and high-dimensional perceptual inputs, with policy synthesis remaining non-parametric and lightweight.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Contribution 3: Extensive validation demonstrating efficiency and performance

Description: The authors validate GPI across diverse simulation benchmarks and real robot platforms, demonstrating that it achieves higher success rates than diffusion-based policies while running substantially faster (20× or more), requiring less memory, and maintaining robustness to perturbations.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Geometry-aware Policy Imitation [View paper](#)
- [1] Imitation learning of neural spatio-temporal point processes [View paper](#)
- [2] Universal planning networks: Learning generalizable representations for visuomotor control [View paper](#)
- [3] A Trajectory Optimisation-Based Incremental Learning Strategy for Learning from Demonstration [View paper](#)
- [4] Diff-Ild: Contact-aware model-based learning from visual demonstration for robotic manipulation via differentiable physics-based simulation and rendering [View paper](#)
- [5] Flow matching-based autonomous driving planning with advanced interactive behavior modeling [View paper](#)
- [6] Peg-in-hole assembly skill imitation learning method based on ProMPs under task geometric representation [View paper](#)
- [7] Adapting a world model for trajectory following in a 3d game [View paper](#)
- [8] Learning Rhythmic Trajectories With Geometric Constraints for Laser-Based Skincare Procedures [View paper](#)
- [9] A structured prediction approach for robot imitation learning [View paper](#)
- [10] Generalizable imitation learning from observation via inferring goal proximity [View paper](#)
- [11] Imitation Learning with Sinkhorn Distances [View paper](#)
- [12] Adversarial Imitation Learning Based on Weighted Wasserstein Distance [View paper](#)
- [13] Sd-pdmd: Deep reinforcement learning for robotic trajectory imitation [View paper](#)
- [14] DFields: Dynamic 3D Descriptor Fields for Zero-Shot Generalizable Robotic Manipulation [View paper](#)
- [15] An approach for imitation learning on Riemannian manifolds [View paper](#)
- [16] One-shot learning for rapid generation of structured robotic manipulation tasks from 3D video demonstrations [View paper](#)
- [17] Complex Robot Motion Generation via Context-Driven Sequencing of Movement Primitives [View paper](#)
- [18] Learning aggressive animal locomotion skills for quadrupedal robots solely from monocular videos [View paper](#)
- [19] Demonstration-Based Proximal Policy Optimization with Action Guidance [View paper](#)
- [20] An imitation from observation approach for dozing distance learning in autonomous bulldozer operation [View paper](#)
- [21] Distance to frontier: From imitation to innovation [View paper](#)
- [22] Neighborly Gains? The Green Spillover Effects of Low-Carbon City Pilot Policy [View paper](#)
- [23] Tracing curves in the plane: Geometric-invariant learning from human demonstrations. [View paper](#)
- [24] Learning to Imitate Spatial Organization in Multi-robot Systems [View paper](#)
- [25] Align Your Intent: Offline Imitation Learning via Optimal Transport [View paper](#)
- [26] Expert Proximity as Surrogate Rewards for Single Demonstration Imitation Learning [View paper](#)
- [27] Problems in the Difference-in-Distance measure of phonetic imitation [View paper](#)
- [28] Cognitive Control Structures in the Imitation Learning of Spatial Sequences and Rhythms-An fMRI Study. [View paper](#)
- [29] Inferring geometric constraints in human demonstrations [View paper](#)
- [30] A Geometric Perspective on Visual Imitation Learning [View paper](#)
- [31] Methods for Teaching Diverse Robot Skills: Leveraging Priors, Geometry, and Dynamics [View paper](#)
- [32] Keypoint Action Tokens Enable In-Context Imitation Learning in Robotics [View paper](#)
- [33] AdaFlow: Imitation Learning with Variance-Adaptive Flow-Based Policies [View paper](#)
- [34] Variational distillation of diffusion policies into mixture of experts [View paper](#)
- [35] Hierarchically branched diffusion models for efficient and interpretable multi-class conditional generation [View paper](#)
- [36] PIDiff: Integrating a High-Performance Transformer Into Diffusion Models for Robust and Efficient Imitation Learning [View paper](#)
- [37] Skill disentanglement for imitation learning from suboptimal demonstrations [View paper](#)
- [38] Goal-conditioned imitation learning using score-based diffusion policies [View paper](#)
- [39] DVANet: Disentangling view and action features for multi-view action recognition [View paper](#)
- [40] Imitation Learning Based on Disentangled Representation Learning of Behavioral Characteristics [View paper](#)

- [41] Discrete policy: Learning disentangled action space for multi-task robotic manipulation [View paper](#)
- [42] Contrasting disentangled partial observations for pedestrian action prediction [View paper](#)
- [43] Modeling human driving behavior through generative adversarial imitation learning [View paper](#)
- [44] Unsupervised learning of disentangled representations from video [View paper](#)
- [45] Disentangled neural relational inference for interpretable motion prediction [View paper](#)
- [46] Latent policy barrier: Learning robust visuomotor policies by staying in-distribution [View paper](#)
- [47] A visual imitation learning algorithm for the selection of robots' grasping points [View paper](#)
- [48] Learning from demonstrations through the use of non-rigid registration [View paper](#)
- [49] Action-Constrained Imitation Learning [View paper](#)
- [50] DexRepNet: Learning Dexterous Robotic Grasping Network with Geometric and Spatial Hand-Object Representations [View paper](#)
- [51] Online Imitation Learning for Manipulation via Decaying Relative Correction through Teleoperation [View paper](#)
- [52] Control framework for collaborative robot using imitation learning-based teleoperation from human digital twin to robot digital twin [View paper](#)
- [53] Motion Tracks: A Unified Representation for Human-Robot Transfer in Few-Shot Imitation Learning [View paper](#)
- [54] Development of an imitation learning method for a neural network system of mobile robot's movement on example of the maze solving [View paper](#)