Novelty Assessment Report

Paper: Goal Reaching with Eikonal-Constrained Hierarchical Quasimetric Reinforcement Learning
PDF URL: https://openreview.net/pdf?id=5WhsCB0Vty
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-07

Abstract

Goal-Conditioned Reinforcement Learning (GCRL) mitigates the difficulty of reward design by framing tasks as goal reaching rather than maximizing hand-crafted reward signals. In this setting, the optimal goal-conditioned value function naturally forms a quasimetric, motivating Quasimetric RL (QRL), which constrains value learning to quasimetric mappings and enforces local consistency through discrete, trajectory-based constraints. We propose Eikonal-Constrained Quasimetric RL (Eik-QRL), a continuous-time reformulation of QRL based on the Eikonal Partial Differential Equation (PDE). This PDE-based structure makes Eik-QRL trajectory-free, requiring only sampled states and goals, while improving out-of-distribution generalization. We provide theoretical guarantees for Eik-QRL and identify limitations that arise under complex dynamics. To address these challenges, we introduce Eik-Hierarchical QRL (Eik-HiQRL), which integrates Eik-QRL into a hierarchical decomposition. Empirically, Eik-HiQRL achieves state-of-the-art performance in offline goal-conditioned navigation and yields consistent gains over QRL in manipulation tasks, matching temporal-difference methods.

Core Task Landscape

This paper addresses: **Goal-Conditioned Reinforcement Learning with Quasimetric Value Functions**

A total of **15 papers** were analyzed and organized into a taxonomy with **14 categories**.

Taxonomy Overview

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

- **Quasimetric Learning Frameworks and Theoretical Foundations**
- **Hierarchical and Compositional Quasimetric Methods**
- **Offline Goal-Conditioned Quasimetric Learning**
- **Asymmetric Cost and Safety-Aware Navigation**
- **Quasimetric Integration with Alternative Paradigms**

Complete Taxonomy Tree

- Goal-Conditioned Reinforcement Learning with Quasimetric Value Functions Survey Taxonomy
- Quasimetric Learning Frameworks and Theoretical Foundations
  - Optimal Quasimetric Value Function Learning (1 papers)
  - [2] Optimal goal-reaching reinforcement learning via quasimetric learning (Wang TongZhou, 2023) View paper
  - Dense Reward Quasimetric Extensions (2 papers)
  - [9] Quasimetric Value Functions with Dense Rewards (Banerjee, 2024) View paper
  - [13] Quasipseudometric Value Functions with Dense Rewards (K Valieva, 2025) View paper
  - Multistep Temporal Distance Estimation (1 papers)
  - [12] Multistep Quasimetric Learning for Scalable Goal-conditioned Reinforcement Learning (Bill Chunyuan Zheng, 2025) View paper
- Hierarchical and Compositional Quasimetric Methods
  - Continuous-Time Eikonal-Based Hierarchical Methods ★ (1 papers)
  - [0] Goal Reaching with Eikonal-Constrained Hierarchical Quasimetric Reinforcement Learning (Anon et al., 2026) View paper
  - Contrastive Learning Hierarchical Integration (1 papers)
  - [11] Hierarchical Quasimetric Reinforcement Learning (Kaiqiang Ke, 2025) View paper
  - Planning Quasi-Metric Decomposition (1 papers)
  - [8] Multi-task reinforcement learning with a planning quasi-metric (Micheli, 2020) View paper
- Offline Goal-Conditioned Quasimetric Learning
  - Contrastive Representation Unification (1 papers)
  - [1] Offline Goal-conditioned Reinforcement Learning with Quasimetric Representations (Myers, 2025) View paper
  - Projective Compositional Planning (1 papers)
  - [3] Offline Goal-Conditioned Reinforcement Learning with Projective Quasimetric Planning (Radji, 2025) View paper
  - Conservative Implicit Value Learning (1 papers)
  - [14] Conservative Offline Goal-Conditioned Implicit V-Learning (K Ke, n.d.) View paper
- Asymmetric Cost and Safety-Aware Navigation
  - Constrained Quasimetric Navigation (1 papers)
  - [4] QuasiNav: Asymmetric Cost-Aware Navigation Planning with Constrained Quasimetric Reinforcement Learning (Hossain, 2025) View paper
  - Quasi-Potential Path-Dependent Cost Decomposition (1 papers)
  - [15] QPRL: Learning Optimal Policies with Quasi-Potential Functions for Asymmetric Traversal (Hossain, n.d.) View paper

- Quasimetric Integration with Alternative Paradigms
  - Decision Transformer Quasimetric Guidance (1 papers)
  - [5] Quasimetric decision transformer: enhancing goal-conditioned reinforcement learning with structured distance guidance (Goyani, 2025) View paper
  - Adversarial Intrinsic Motivation with Quasimetrics (1 papers)
  - [10] Adversarial Intrinsic Motivation for Reinforcement Learning (Ishan Durugkar, 2022) View paper
  - State Representation and Curriculum Learning (2 papers)
  - [6] State chrono representation for enhancing generalization in reinforcement learning (Jianda Chen, 2024) View paper
  - [7] Diffusion-based curriculum reinforcement learning (Iacca, 2024) View paper

Narrative

Core task: goal-conditioned reinforcement learning with quasimetric value functions. The field centers on learning asymmetric distance-like functions that capture the cost or difficulty of reaching one state from another, enabling agents to plan toward arbitrary goals. The taxonomy reveals five main branches: foundational quasimetric learning frameworks that establish theoretical properties and basic algorithms; hierarchical and compositional methods that exploit quasimetric structure for temporal abstraction and subgoal decomposition; offline approaches that learn quasimetrics from fixed datasets without environment interaction; asymmetric cost and safety-aware navigation techniques that handle directional constraints or obstacle avoidance; and integration with alternative paradigms such as diffusion models or transformers. Representative works like Optimal Quasimetric Learning[2] and Planning Quasi-Metric[8] anchor the foundational branch, while Quasimetric Decision Transformer[5] and Offline Quasimetric Representations[1] illustrate offline and transformer-based extensions.

Recent activity highlights contrasts between online hierarchical methods and offline data-driven approaches, as well as trade-offs between theoretical rigor and practical scalability. Hierarchical methods such as Hierarchical Quasimetric[11] and Multistep Quasimetric[12] decompose long-horizon tasks into subgoals, yet face challenges in continuous-time settings where smooth value propagation is critical. The original paper, Eikonal Hierarchical Quasimetric[0], sits within the hierarchical and compositional branch, specifically addressing continuous-time dynamics through Eikonal-based formulations. Compared to discrete hierarchical works like Hierarchical Quasimetric[11], it emphasizes differential equations for value function smoothness, while differing from offline methods such as Offline Quasimetric Representations[1] by focusing on online learning with temporal abstraction. Open questions remain around balancing asymmetry with computational efficiency and integrating safety constraints into hierarchical quasimetric planning.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

The original leaf focuses on continuous-time, trajectory-free hierarchical RL using Eikonal PDE formulations for quasimetric value functions. The sibling subtopics represent alternative hierarchical approaches: one integrating contrastive learning for goal space reduction, and another decomposing planning into environment models and task-specific components. All three share hierarchical quasimetric RL as a foundation but differ fundamentally in their mathematical frameworks and decomposition strategies.

Similarities: - All three subtopics operate within hierarchical goal-conditioned RL frameworks using quasimetric value functions - Each approach aims to improve scalability or efficiency in multi-goal or multi-task settings - All exclude standard discrete trajectory-based methods, indicating a focus on more sophisticated temporal or spatial abstractions

Differences: - The original leaf uses continuous-time Eikonal PDE formulations and is trajectory-free, while siblings use discrete learning paradigms - Contrastive Learning Hierarchical Integration explicitly incorporates contrastive learning for goal space construction, which the original leaf excludes - Planning Quasi-Metric Decomposition separates task-agnostic environment modeling from task-specific aiming, while the original leaf maintains a unified Eikonal-based formulation - The original leaf's PDE-based approach is mathematically distinct from the contrastive and planning decomposition methods used by siblings

Suggested Search Directions: - Hybrid methods combining Eikonal PDEs with contrastive learning or planning decompositions - Comparative studies on continuous-time vs discrete hierarchical quasimetric approaches - Applications of different hierarchical quasimetric methods to specific domains (robotics, navigation, manipulation)

Sibling Subtopics

- **Contrastive Learning Hierarchical Integration** (leaves: 1, papers: 1)
- Scope: Hierarchical frameworks combining quasimetric learning with contrastive learning for reduced goal space construction.
- Exclude: Non-hierarchical contrastive methods, Eikonal-based approaches, and planning decompositions belong elsewhere.
- **Planning Quasi-Metric Decomposition** (leaves: 1, papers: 1)
- Scope: Multi-task methods decomposing quasimetric planning into task-agnostic environment models and task-specific aimers.
- Exclude: Single-task methods, hierarchical value decompositions, and offline-specific approaches belong in other categories.

Contributions Analysis

Overall novelty summary. The paper introduces Eikonal-Constrained Quasimetric RL (Eik-QRL) and its hierarchical extension Eik-HiQRL, reformulating quasimetric learning through continuous-time Eikonal PDEs rather than discrete trajectory constraints. Within the taxonomy, it occupies the 'Continuous-Time Eikonal-Based Hierarchical Methods' leaf under the hierarchical branch. Notably, this leaf contains only the original paper itself—no sibling papers appear in the same category. This isolation suggests the continuous-time Eikonal formulation for hierarchical quasimetric learning represents a relatively unexplored direction within the field's current landscape.

The taxonomy reveals neighboring leaves focused on contrastive learning integration and planning-based decomposition within the hierarchical branch, plus foundational quasimetric methods and offline approaches in adjacent branches. The scope note for the original paper's leaf explicitly excludes 'discrete trajectory-based methods' and 'contrastive learning integrations,' positioning Eik-QRL as distinct from both the foundational discrete methods and alternative hierarchical strategies. The broader hierarchical branch contains only three leaves total, indicating that hierarchical quasimetric methods remain a moderately sparse research direction compared to the foundational and offline branches.

Among four candidates examined across three contributions, no refutable prior work was identified. The Eik-HiQRL contribution examined three candidates with none providing clear overlap, while the theoretical guarantees examined one candidate without refutation. The core Eik-QRL contribution examined zero candidates, likely reflecting the novelty of the continuous-time PDE formulation. Given the limited search scope—only four total candidates across all contributions—these statistics suggest the specific combination of Eikonal constraints and hierarchical decomposition has minimal direct precedent among semantically similar papers, though the small sample size precludes definitive conclusions about the broader literature.

Based on top-K semantic search examining four candidates, the work appears to occupy a sparse intersection of continuous-time PDE methods and hierarchical quasimetric learning. The taxonomy structure confirms limited prior activity in this specific direction, with the original paper as the sole occupant of its leaf. However, the restricted search scope means potentially relevant work in adjacent areas—

such as PDE-based RL outside the quasimetric framework or hierarchical methods using alternative formulations—may not have been captured in this analysis.

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

Contribution 1: Eikonal-Constrained Quasimetric RL (Eik-QRL)

Description: The authors introduce Eik-QRL, a novel formulation that reformulates Quasimetric RL using continuous-time constraints derived from the Eikonal PDE rather than discrete trajectory-based constraints. This PDE-based structure makes the approach trajectory-free, requiring only sampled states and goals, and improves out-of-distribution generalization.

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: Eikonal-Hierarchical QRL (Eik-HiQRL)

Description: The authors propose Eik-HiQRL, a hierarchical algorithm that addresses the limitations of Eik-QRL under complex dynamics by integrating Eik-QRL into a hierarchical framework. This design combines accurate quasimetric projection in low-dimensional abstract spaces with PDE-based advantages and hierarchical structure to improve signal-to-noise ratio in long-horizon tasks.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Transferring task goals via hierarchical reinforcement learning

URL: View paper

Brief Assessment

Transferring Hierarchical Goals[17] focuses on transferring high-level task understanding across different robot bodies in navigation tasks, using subgoals to separate 'what to do' from 'how to do it'. The original paper's Eik-HiQRL addresses a different problem: integrating Eikonal PDE constraints into hierarchical quasimetric learning for long-horizon goal-conditioned RL, with theoretical guarantees based on continuous-time formulations. These are distinct technical contributions addressing different aspects of hierarchical RL.

2. Graph-Assisted Stitching for Offline Hierarchical Reinforcement Learning

URL: View paper

Brief Assessment

Graph-Assisted Stitching[16] focuses on graph-based subgoal selection without learning a high-level policy, using temporal distance representations for clustering. This differs from Eik-HiQRL's PDE-constrained quasimetric learning in abstract spaces with hierarchical value decomposition.

3. Hierarchical Quasimetric Reinforcement Learning

URL: View paper

Brief Assessment

Hierarchical Quasimetric[11] uses contrastive learning and imitation learning in its hierarchical framework, whereas Eik-HiQRL integrates Eik-QRL (PDE-based quasimetric learning) with hierarchical decomposition for improved signal-to-noise ratio in long-horizon tasks.

Contribution 3: Theoretical guarantees for Eik-QRL

Description: The authors establish theoretical guarantees for Eik-QRL, including optimal value recovery under regularity conditions (Lemma 4.7 and Theorem 4.8), and identify inherent limitations when the method is applied to complex dynamical settings. This analysis provides formal justification for the hierarchical extension.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. QuasiNav: Asymmetric Cost-Aware Navigation Planning with Constrained Quasimetric Reinforcement Learning

URL: View paper

Brief Assessment

QuasiNav[4] focuses on navigation planning with quasimetric embeddings for asymmetric traversal costs in outdoor environments, not on theoretical guarantees for quasimetric RL under regularity conditions and complex dynamics as in the original paper.

Appendix: Text Similarity Detection

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

References

- [0] Goal Reaching with Eikonal-Constrained Hierarchical Quasimetric Reinforcement Learning View paper
- [1] Offline Goal-conditioned Reinforcement Learning with Quasimetric Representations View paper
- [2] Optimal goal-reaching reinforcement learning via quasimetric learning View paper
- [3] Offline Goal-Conditioned Reinforcement Learning with Projective Quasimetric Planning View paper
- [4] QuasiNav: Asymmetric Cost-Aware Navigation Planning with Constrained Quasimetric Reinforcement Learning View paper
- [5] Quasimetric decision transformer: enhancing goal-conditioned reinforcement learning with structured distance guidance View paper
- [6] State chrono representation for enhancing generalization in reinforcement learning View paper
- [7] Diffusion-based curriculum reinforcement learning View paper
- [8] Multi-task reinforcement learning with a planning quasi-metric View paper
- [9] Quasimetric Value Functions with Dense Rewards View paper
- [10] Adversarial Intrinsic Motivation for Reinforcement Learning View paper
- [11] Hierarchical Quasimetric Reinforcement Learning View paper
- [12] Multistep Quasimetric Learning for Scalable Goal-conditioned Reinforcement Learning View paper
- [13] Quasipseudometric Value Functions with Dense Rewards View paper
- [14] Conservative Offline Goal-Conditioned Implicit V-Learning View paper

- [15] QPRL: Learning Optimal Policies with Quasi-Potential Functions for Asymmetric Traversal View paper
- [16] Graph-Assisted Stitching for Offline Hierarchical Reinforcement Learning View paper
- [17] Transferring task goals via hierarchical reinforcement learning View paper