Novelty Assessment Report

Paper: Discovering Diverse Behaviors via Temporal Contrastive Learning

PDF URL: https://openreview.net/pdf?id=KjYpHySlb0

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Effective exploration in reinforcement learning requires not only tracking where an agent has been, but also understanding how the agent perceives and represents the world. To learn powerful representations, an agent should actively explore states that contribute to its knowledge of the environment. Temporal representations can capture the information necessary to solve a wide range of potential tasks while avoiding the computational cost associated with full state reconstruction. In this paper, we propose an exploration method that leverages temporal contrastive representations to guide exploration, prioritizing states with unpredictable future outcomes. We demonstrate that such representations can enable the learning of complex exploratory behaviors in locomotion, manipulation, and embodied-AI tasks, revealing capabilities and behaviors that traditionally require extrinsic rewards. Unlike approaches that rely on explicit distance learning or episodic memory mechanisms (e.g., quasimetric-based methods), our method builds directly on temporal similarities, yielding a simpler yet effective strategy for exploration.

Core Task Landscape

This paper addresses: **Exploration via Temporal Contrastive Representations**

A total of **27 papers** were analyzed and organized into a taxonomy with **12 categories**.

Taxonomy Overview

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

- **Temporal Contrastive Representation Learning for Exploration and Control**
- **Temporal Contrastive Learning for Transfer and Multi-Agent Coordination**
- **Temporal Contrastive Learning for Time-Series and Non-RL Applications**
- **Temporal Action Representation Learning for Resource-Aware Decision-Making**
- **Generative AI and Temporal Contrastive Learning in RL**

Complete Taxonomy Tree

- Exploration via Temporal Contrastive Representations Survey Taxonomy
- Temporal Contrastive Representation Learning for Exploration and Control
  - Curiosity-Driven Exploration via Temporal Contrastive Representations ★ (4 papers)
  - [0] Discovering Diverse Behaviors via Temporal Contrastive Learning (Anon et al., 2026) View paper
  - [11] Episodic novelty through temporal distance (Jiang Yu-hua, 2025) View paper
  - [17] Self-Supervised Exploration via Temporal Inconsistency in Reinforcement Learning (Zijian Gao, 2024) View paper
  - [24] Curiosity-Driven Exploration via Temporal Contrastive Learning (F Mohamed, n.d.) View paper
  - Temporal Contrastive Learning for Visual Control and Visuomotor Policies (4 papers)
  - [5] : Temporal Latent Action-Driven Contrastive Loss for Visual Reinforcement Learning (R Zheng, 2023) View paper
  - [9] Premier-TACO is a Few-Shot Policy Learner: Pretraining Multitask Representation via Temporal Action-Driven Contrastive Loss (Zheng, 2024) View paper
  - [20] Oracle-Guided Masked Contrastive Reinforcement Learning for Visuomotor Policies (Zhang Yuhang, 2025) View paper
  - [22] CST-RL: Contrastive Spatio-Temporal Representations for Reinforcement Learning (Chi-Kai Ho, 2023) View paper
  - Spatial-Temporal Contrastive Representation Learning (3 papers)
  - [3] STACoRe: Spatio-temporal and action-based contrastive representations for reinforcement learning in Atari (Young-Jae Lee, 2023) View paper
  - [4] BEVNav: Robot Autonomous Navigation via Spatial-Temporal Contrastive Learning in Bird's-Eye View (Jiahao Jiang, 2024) View paper
  - [6] MOOSS: Mask-Enhanced Temporal Contrastive Learning for Smooth State Evolution in Visual Reinforcement Learning (Jiarui Sun, 2024) View paper
  - Temporal Contrastive Learning for Hierarchical and Goal-Conditioned RL (3 papers)
  - [2] Inference via Interpolation: Contrastive Representations Provably Enable Planning and Inference (Eysenbach, 2024) View paper
  - [7] Temporal abstractions-augmented temporally contrastive learning: An alternative to the Laplacian in RL (Erraqabi, 2022) View paper
  - [13] Balancing Exploration and Exploitation in Hierarchical Reinforcement Learning via Latent Landmark Graphs (Qingyang Zhang, 2023) View paper
- Temporal Contrastive Learning for Transfer and Multi-Agent Coordination
  - Transfer Learning via Temporal Contrastive Representations (2 papers)
  - [18] Multi-Agent Transfer Learning via Temporal Contrastive Learning (Zeng, 2024) View paper
  - [19] Transfer Learning via Temporal Contrastive Learning (Zeng, 2024) View paper
  - Inverse Reinforcement Learning with Temporal Contrastive Representations (1 papers)

- [1] TW-CRL: Time-Weighted Contrastive Reward Learning for Efficient Inverse Reinforcement Learning (Li, 2025) View paper

- Temporal Contrastive Learning for Time-Series and Non-RL Applications
  - Self-Supervised Temporal Contrastive Learning for Time-Series Representation (4 papers)
  - [8] Self-supervised group meiosis contrastive learning for EEG-based emotion recognition: H. Kan et al. (H Kan, 2023) View paper
  - [14] Timedrl: Disentangled representation learning for multivariate time-series (Ching Chang, 2024) View paper
  - [15] Finding Order in Chaos: A Novel Data Augmentation Method for Time Series in Contrastive Learning (Demirel, 2023) View paper
  - [23] Contrastive-Learning-Based Time-Series Feature Representation for Parcel-Based Crop Mapping Using Incomplete Sentinel-2 Image Sequences (Yanan Zhou, 2023) View paper
  - Temporal Contrastive Learning for Anomaly Detection and Active Learning (2 papers)
  - [12] Carla: A self-supervised contrastive representation learning approach for time series anomaly detection (ZZ Darban, 2023) View paper
  - [25] Contrastive Representation based Active Learning for Time Series (Lujia Pan, 2022) View paper
  - Temporal Contrastive Learning for Imitation and Video-Based Representation (2 papers)
  - [10] Towards principled representation learning from videos for reinforcement learning (Misra, 2024) View paper
  - [16] Time-contrastive networks: Self-supervised learning from video (Sermanet, 2018) View paper
  - Cross-Modal Temporal Contrastive Learning (1 papers)
  - [21] What Time Tells Us? An Explorative Study of Time Awareness Learned from Static Images (Hu Han, 2025) View paper
- Temporal Action Representation Learning for Resource-Aware Decision-Making (1 papers)
  - [26] Temporal Action Representation Learning for Aerial Maneuvering and Resource-Aware Decision-Making (H Jung, n.d.) View paper
- Generative AI and Temporal Contrastive Learning in RL (1 papers)
  - [27] Notes from an ICML 2025 Tutorial: Generative AI Meets Reinforcement Learning (B Eysenbach, n.d.) View paper

Narrative

Core task: exploration via temporal contrastive representations. This field leverages contrastive learning over temporal sequences to build representations that guide exploration and decision-making in reinforcement learning and beyond. The taxonomy reveals five main branches. The first branch focuses on temporal contrastive representation learning for exploration and control, encompassing curiosity-driven methods that use temporal contrasts to identify novel states or behaviors. A second branch examines transfer and multi-agent coordination, where temporal contrastive objectives help agents generalize across tasks or coordinate in shared environments. A third branch extends these ideas to time-series analysis and non-RL applications such as video understanding and crop mapping. The fourth branch targets resource-aware decision-making, using temporal action representations to optimize computational or physical constraints. Finally, a fifth branch explores the intersection of generative AI and temporal contrastive learning in RL, investigating how generative models can enhance or complement contrastive objectives.

Within the exploration and control branch, many studies design intrinsic rewards or curiosity signals by contrasting temporally adjacent or distant observations. Temporal Contrastive Behaviors[0] sits squarely in this curiosity-driven cluster, alongside works like Curiosity Temporal Contrastive[24] and Temporal Inconsistency Exploration[17], which similarly exploit temporal structure to drive exploration. Compared to Episodic Novelty Distance[11], which measures novelty via episodic memory, Temporal Contrastive Behaviors[0] emphasizes learning contrastive embeddings that capture behavioral dynamics over time. Meanwhile, foundational methods such as Time-Contrastive Networks[16] established early principles for temporal contrast, and recent efforts like STACoRe[3] and Premier-TACO[9] refine these ideas with state-action or hierarchical abstractions. The central trade-off across these lines involves balancing the granularity of temporal contrasts—whether to compare consecutive frames, longer horizons, or abstract action sequences—against computational cost and sample efficiency.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Episodic novelty through temporal distance

Authors: Jiang Yu-hua, Liu, Qihan, Yuhua Jiang, Yang, et al. (25 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Exploration in sparse reward environments remains a significant challenge in reinforcement learning, particularly in Contextual Markov Decision Processes (CMDPs), where environments differ across episodes. Existing episodic intrinsic motivation methods for CMDPs primarily rely on count-based approaches, which are ineffective in large state spaces, or on similarity-based methods that lack appropriate metrics for state comparison. To address these shortcomings, we propose Episodic Novelty Through ...

Relationship Analysis

Both papers belong to the curiosity-driven exploration via temporal contrastive representations category, using temporal contrastive learning to generate intrinsic rewards for exploration. They overlap in leveraging contrastive learning to capture temporal structure and guide exploration based on state novelty or predictability. The key differences are: the original paper (C-TeC) rewards states with unpredictable futures using direct temporal similarity without episodic memory or quasimetric constraints, while the candidate paper (ETD) explicitly learns temporal distances via quasimetric parameterization and uses episodic memory to compute minimum distances from past states as intrinsic rewards.

---

2. Self-Supervised Exploration via Temporal Inconsistency in Reinforcement Learning

Authors: Zijian Gao, Kele Xu, Yuanzhao Zhai, Bo Ding, Dawei Feng, et al. (8 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

In sparse extrinsic reward settings, reinforcement learning remains a challenge despite increasing interest in this field. Existing approaches suggest that intrinsic rewards can alleviate issues caused by reward sparsity. However, many studies overlook the critical role of temporal information, essential for human curiosity. This article introduces a novel intrinsic reward mechanism inspired by human learning processes, where curiosity is evaluated by comparing current observations with historic...

Relationship Analysis

Both papers belong to the curiosity-driven exploration via temporal contrastive representations category, using temporal structure to generate intrinsic rewards for exploration. They overlap in leveraging temporal prediction models to guide exploration without extrinsic rewards, but differ fundamentally in their approach: the original paper uses temporal contrastive learning with InfoNCE loss to measure future state predictability and rewards unpredictable futures, while the candidate paper uses temporal inconsistency across multiple snapshots of a prediction model measured via weighted nuclear norm to assess exploration novelty.

---

3. Curiosity-Driven Exploration via Temporal Contrastive Learning

Authors: F Mohamed, C Ji, B Eysenbach, G Berseth | **URL**: View paper

Abstract

â¦ To improve exploration, we learn representations that encode the agentâs future state â¦ using temporal contrastive learning. We begin by describing how contrastive representation â¦

⚠ Similarity Notice

These papers appear to be highly similar or the same work, both titled with variations of 'Curiosity-Driven Exploration via Temporal Contrastive Learning' and describing identical core methods: using temporal contrastive representations to generate intrinsic rewards by rewarding states with unpredictable futures, avoiding episodic memory and quasimetric learning. The technical approach, experimental setup, and core contributions are nearly identical, suggesting these are likely different versions or submissions of the same paper.

Contributions Analysis

Overall novelty summary. The paper proposes an exploration method that uses temporal contrastive representations to prioritize states with unpredictable future outcomes, generating intrinsic rewards from prediction errors in learned embeddings. It resides in the 'Curiosity-Driven Exploration via Temporal Contrastive Representations' leaf, which contains four papers total (including this one). This leaf sits within a broader branch of 'Temporal Contrastive Representation Learning for Exploration and Control' (four leaves, approximately 16 papers), indicating a moderately active research direction. The taxonomy shows this is not an isolated niche but part of a structured exploration-focused subfield.

The taxonomy reveals neighboring leaves focused on visual control (four papers), spatial-temporal fusion (three papers), and hierarchical goal-conditioned RL (three papers). These adjacent directions share the use of temporal contrastive objectives but diverge in application: visual control emphasizes visuomotor policies from high-dimensional observations, while hierarchical methods discover subgoals for planning. The paper's focus on curiosity-driven exploration distinguishes it from these control-centric approaches, though the underlying contrastive mechanism connects across boundaries. The taxonomy's scope and exclude notes clarify that this work targets intrinsic motivation rather than policy learning or multi-agent coordination.

Among 18 candidates examined, the contribution-level analysis shows mixed novelty signals. The core exploration method (4 candidates examined, 1 refutable) and intrinsic reward mechanism (10 candidates examined, 1 refutable) both encounter at least one overlapping prior work within the limited search scope. The third contribution—positioning as a simpler alternative to quasimetric methods—shows no refutable candidates among 4 examined, suggesting this framing may be more distinctive. The statistics indicate a focused but not exhaustive search: the presence of refutable candidates does not imply the ideas are well-trodden, only that some relevant prior work exists within the top-K semantic matches.

Based on the limited search scope (18 candidates from semantic retrieval), the work appears to build on established principles of temporal contrastive learning for exploration, with some contributions showing overlap in the examined literature. The taxonomy context suggests the paper operates in a moderately populated research area where temporal contrastive methods for curiosity are actively studied. The analysis does not cover exhaustive citation networks or domain-specific venues, so additional related work may exist beyond the examined candidates.

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

Contribution 1: Exploration method using temporal contrastive representations

Description: The authors introduce an exploration approach that uses temporal contrastive learning to learn representations capturing future state occupancy. The method rewards agents for visiting states with unpredictable futures, enabling complex exploratory behaviors without extrinsic rewards.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Temporal abstractions-augmented temporally contrastive learning: An alternative to the Laplacian in RL

URL: View paper

Brief Assessment

Temporal Abstractions Contrastive[7] focuses on learning Laplacian representations in non-uniform-prior settings using skill-based covering policies, not on exploration via unpredictable future states as the primary objective.

---

2. Curiosity-Driven Exploration via Temporal Contrastive Learning

URL: View paper

Prior Art Analysis

Curiosity Temporal Contrastive[24] demonstrates that the core approach of using temporal contrastive learning for exploration with rewards based on unpredictable futures was previously published. The candidate paper presents the same fundamental method: learning temporal contrastive representations to capture future state occupancy, then rewarding agents for visiting states with unpredictable futures. Both papers use identical mathematical formulations (discounted state occupancy measures, InfoNCE loss, negative similarity scores as rewards), the same algorithmic structure, and evaluate on overlapping environments. The candidate paper explicitly describes this as their main contribution and provides the complete technical framework that the original paper claims as novel.

Evidence

Evidence 1 - **Rationale**: Both papers propose the same core exploration method using temporal contrastive representations. The candidate explicitly states this as their contribution, demonstrating prior work exists. - **Original**: we propose an exploration method that leverages temporal contrastive representations to guide exploration, prioritizing states with unpredictable future outcomes. - **Candidate**: we propose an exploration method that leverages temporal contrastive representations to guide exploration, aiming to maximize state coverage as perceived through the lens of these learned representations.

Evidence 2 - **Rationale**: Both papers describe the identical mechanism for extracting exploration signals: negating the contrastive similarity score to reward unpredictable futures. This is the core technical contribution claimed as novel. - **Original**: the contrastive model produces a similarity score between state-action pairs (st, at) and future states sf . negating this similarity score results in our exploration signal rintr, encouraging the agent to visit states that are not predictive of future states in the same trajectory in the eyes of th... - **Candidate**: the contrastive model produces a similarity score between state-action pairs (st, at) and future states sf proportional to the probability of reaching sf from (st, at). negating this similarity score results in our exploration signal rintr, encouraging the agent to visit states that appear to have h...

Evidence 3 - **Rationale**: The mathematical formulation of the intrinsic reward is identical in both papers, showing the candidate published this exact technical approach first. - **Original**: e[rintr(st, at)] = ept (sf |st,at) [-cθ((st, at), sf)] = ept (sf |st,at) [||ϕθ(st, at) - ψθ(sf)||] - **Candidate**: e[rintr(st, at)] = ept (sf |st,at) [-cθ((st, at), sf)] = ept (sf |st,at) [||ϕθ(st, at) - ψθ(sf)||]

Evidence 4 - **Rationale**: Both papers use identical visual explanations and descriptions of the core mechanism: learning temporal similarity and rewarding unpredictable futures. - **Original**: we learn temporal representations so that the representation of (s0, a0) is more similar to (s2,3,4,...). we reward the agent for visiting future states that seem unpredictable. - **Candidate**: we train a contrastive model such that the temporal similarity between the representation of (s0, a0) and (s2,3,4,...) should be high. we reward the agent for visiting states whose futures seem far away/improbable.

Evidence 5 - **Rationale**: Both papers explicitly state their main contribution is a new exploration method based on temporal contrastive representations, with the candidate demonstrating this was published earlier. - **Original**: the main contribution of this work is a new objective for exploration based on the prediction error of temporal representations. we demonstrate our approach by maximizing these intrinsic rewards with ppo (schulman et al., 2017) and sac (haarnoja et al., 2018c). - **Candidate**: the main contribution of this work is a new exploration algorithm that achieves state-of-the-art state coverage across navigation, manipulation, and open-world environments. connections are made between this contrastive learning-based objective and information control objectives

3. Curiosity-driven learning in artificial intelligence and its applications

URL: View paper

Brief Assessment

Curiosity-Driven Learning[39] is a PhD thesis providing a broad survey of curiosity-driven learning across multiple domains (RL, recommender systems, classification). It does not present a specific exploration method using temporal contrastive representations for unpredictable future states as described in the original paper.

4. Contrastive difference predictive coding

URL: View paper

Brief Assessment

Contrastive Difference Predictive[38] focuses on goal-conditioned RL and policy evaluation using temporal difference learning for contrastive predictive coding, not on exploration driven by unpredictable future states as an intrinsic reward signal.

Contribution 2: Intrinsic reward based on temporal representation prediction error

Description: The paper presents a novel intrinsic reward signal derived from the prediction error of learned temporal contrastive representations. This reward encourages agents to explore states that are less informative about future states according to the contrastive model.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Latent world models for intrinsically motivated exploration

URL: View paper

Prior Art Analysis

Latent World Models[31] demonstrates prior work that uses prediction error from temporal representations for intrinsic motivation. The candidate paper presents a method where a world model operating on temporally-arranged latent representations produces prediction errors used as intrinsic rewards for exploration. The candidate explicitly states that 'the novelty and missing information are estimated as a prediction error of the world model' and that this prediction error is normalized and used as an intrinsic reward signal. The temporal representations are learned to arrange embeddings 'respecting temporal distance of observations' through a self-supervised method. This directly parallels the original paper's contribution of deriving intrinsic rewards from prediction errors of learned temporal contrastive representations.

Evidence

Evidence 1 - **Rationale**: Both methods use temporal representations for exploration. The candidate arranges embeddings by temporal distance and uses prediction error for novelty detection, while the original uses temporal contrastive representations to prioritize unpredictable states. - **Original**: we propose an exploration method that leverages temporal contrastive representations to guide exploration, prioritizing states with unpredictable future outcomes - **Candidate**: we present a self-supervised representation learning method for image-based observations, which arranges embeddings respecting temporal distance of observations. this representation is empirically robust to stochasticity and suitable for novelty detection from the error of a predictive forward model...

2. Learning to Explore in Diverse Reward Settings via Temporal-Difference-Error Maximization

URL: View paper

Brief Assessment

TD-Error Maximization[28] uses temporal-difference error of the exploitation objective as exploration reward, not prediction error of learned temporal contrastive representations. The candidate focuses on maximizing absolute TD-error from value function approximation, while the original paper derives intrinsic rewards from contrastive representation learning prediction errors.

3. Temporal difference uncertainties as a signal for exploration

URL: View paper

Brief Assessment

TD Uncertainties Exploration[33] focuses on uncertainty estimation over temporal difference errors in value-based RL, not on temporal contrastive representations for exploration. The candidate uses TD-error variance as an intrinsic reward, while the original paper uses prediction error of learned temporal contrastive representations to guide exploration.

4. Variational state encoding as intrinsic motivation in reinforcement learning

URL: View paper

Brief Assessment

Variational State Encoding[32] uses KL divergence between VAE posterior and prior distributions as intrinsic motivation, not temporal representation prediction error. The candidate focuses on encoding latent structure via VAE rather than predicting future states through temporal contrastive learning.

5. A reinforcement learning method of solving Markov decision processes: an adaptive exploration model based on temporal difference error

URL: View paper

Brief Assessment

Adaptive Exploration Model[29] focuses on temporal difference error between dual networks for value function approximation in MDPs, not on intrinsic rewards derived from temporal contrastive representation prediction errors for exploration.

6. Surprise signals in the supplementary eye field: rectified prediction errors drive exploration-exploitation transitions

URL: View paper

Brief Assessment

Surprise Signals Exploration[36] focuses on neural mechanisms in the supplementary eye field for exploration-exploitation transitions in visual search tasks, not on intrinsic reward signals derived from temporal contrastive representations for reinforcement learning agents.

---

7. Phasic dopamine as a prediction error of intrinsic and extrinsic reinforcements driving both action acquisition and reward maximization: A simulated robotic study

URL: View paper

Brief Assessment

Phasic Dopamine Prediction[35] focuses on biological dopamine mechanisms and intrinsic reinforcements from unexpected environmental events, not on temporal contrastive representations for RL exploration as in the original paper.

---

8. In search of the neural circuits of intrinsic motivation

URL: View paper

Brief Assessment

Neural Circuits Intrinsic[34] focuses on intrinsic motivation through prediction error in general neural circuits, not specifically temporal contrastive representations for RL exploration as in the original paper.

---

9. Tracking emotions: intrinsic motivation grounded on multi-level prediction error dynamics

URL: View paper

Brief Assessment

Multi-Level Prediction Error[30] focuses on prediction error dynamics across multiple hierarchical levels (goal-level and system-level) for robotic visuo-motor coordination, not on temporal contrastive representations for RL exploration as in the original paper.

---

10. Predication-Error-Based Intrinsically Motivated Saccade Learning

URL: View paper

Brief Assessment

Intrinsically Motivated Saccade[37] focuses on saccadic eye movement learning using sensory prediction errors in a specific motor control context, not general RL exploration via temporal contrastive representations across diverse tasks.

Contribution 3: Simpler alternative to quasimetric-based exploration methods

Description: The authors develop a method that avoids the complexity of quasimetric learning and episodic memory used in prior work like ETD. Their approach works directly with temporal similarities, making it more amenable to off-policy RL algorithms while maintaining competitive performance.

This contribution was assessed against **4 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Multistep Quasimetric Learning for Scalable Goal-conditioned Reinforcement Learning

URL: View paper

Brief Assessment

Multistep Quasimetric Learning[40] focuses on goal-conditioned RL using quasimetric architectures for temporal distance estimation, not on exploration methods that avoid quasimetric learning as claimed in the original paper.

---

2. Goal Reaching with Eikonal-Constrained Hierarchical Quasimetric Reinforcement Learning

URL: View paper

Brief Assessment

Eikonal Hierarchical Quasimetric[42] focuses on goal-conditioned RL with quasimetric value functions for goal reaching, not on exploration methods. The original paper addresses exploration via temporal similarity without quasimetric learning or episodic memory, while the candidate uses quasimetric structures for value learning in goal-reaching tasks.

---

3. Quasimetric decision transformer: enhancing goal-conditioned reinforcement learning with structured distance guidance

URL: View paper

Brief Assessment

Quasimetric Decision Transformer[43] focuses on goal-conditioned offline RL using quasimetric distances to replace returns-to-go in decision transformers, not on exploration methods or temporal similarity-based intrinsic rewards for online RL.

---

4. Offline Goal-conditioned Reinforcement Learning with Quasimetric Representations

URL: View paper

Brief Assessment

Offline Quasimetric Representations[41] focuses on offline goal-conditioned RL with quasimetric representations for optimal policy learning, not on simplifying exploration methods. The candidate addresses a different problem domain (goal-reaching from offline data) than the original's exploration without extrinsic rewards.

Appendix: Text Similarity Detection

Textual similarity detection checked 20 papers and found 3 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. Curiosity-Driven Exploration via Temporal Contrastive Learning

Detected in: Core Task (sibling), Contribution: contribution_1

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Discovering Diverse Behaviors via Temporal Contrastive Learning View paper
- [1] TW-CRL: Time-Weighted Contrastive Reward Learning for Efficient Inverse Reinforcement Learning View paper
- [2] Inference via Interpolation: Contrastive Representations Provably Enable Planning and Inference View paper
- [3] STACoRe: Spatio-temporal and action-based contrastive representations for reinforcement learning in Atari View paper
- [4] BEVNav: Robot Autonomous Navigation via Spatial-Temporal Contrastive Learning in Bird's-Eye View View paper
- [5] : Temporal Latent Action-Driven Contrastive Loss for Visual Reinforcement Learning View paper
- [6] MOOSS: Mask-Enhanced Temporal Contrastive Learning for Smooth State Evolution in Visual Reinforcement Learning View paper
- [7] Temporal abstractions-augmented temporally contrastive learning: An alternative to the Laplacian in RL View paper
- [8] Self-supervised group meiosis contrastive learning for EEG-based emotion recognition: H. Kan et al. View paper
- [9] Premier-TACO is a Few-Shot Policy Learner: Pretraining Multitask Representation via Temporal Action-Driven Contrastive Loss View paper
- [10] Towards principled representation learning from videos for reinforcement learning View paper
- [11] Episodic novelty through temporal distance View paper
- [12] Carla: A self-supervised contrastive representation learning approach for time series anomaly detection View paper
- [13] Balancing Exploration and Exploitation in Hierarchical Reinforcement Learning via Latent Landmark Graphs View paper
- [14] Timedrl: Disentangled representation learning for multivariate time-series View paper
- [15] Finding Order in Chaos: A Novel Data Augmentation Method for Time Series in Contrastive Learning View paper
- [16] Time-contrastive networks: Self-supervised learning from video View paper
- [17] Self-Supervised Exploration via Temporal Inconsistency in Reinforcement Learning View paper
- [18] Multi-Agent Transfer Learning via Temporal Contrastive Learning View paper
- [19] Transfer Learning via Temporal Contrastive Learning View paper
- [20] Oracle-Guided Masked Contrastive Reinforcement Learning for Visuomotor Policies View paper
- [21] What Time Tells Us? An Explorative Study of Time Awareness Learned from Static Images View paper
- [22] CST-RL: Contrastive Spatio-Temporal Representations for Reinforcement Learning View paper
- [23] Contrastive-Learning-Based Time-Series Feature Representation for Parcel-Based Crop Mapping Using Incomplete Sentinel-2 Image Sequences View paper
- [24] Curiosity-Driven Exploration via Temporal Contrastive Learning View paper
- [25] Contrastive Representation based Active Learning for Time Series View paper
- [26] Temporal Action Representation Learning for Aerial Maneuvering and Resource-Aware Decision-Making View paper
- [27] Notes from an ICML 2025 Tutorial: Generative AI Meets Reinforcement Learning View paper
- [28] Learning to Explore in Diverse Reward Settings via Temporal-Difference-Error Maximization View paper
- [29] A reinforcement learning method of solving Markov decision processes: an adaptive exploration model based on temporal difference error View paper
- [30] Tracking emotions: intrinsic motivation grounded on multi-level prediction error dynamics View paper
- [31] Latent world models for intrinsically motivated exploration View paper
- [32] Variational state encoding as intrinsic motivation in reinforcement learning View paper
- [33] Temporal difference uncertainties as a signal for exploration View paper
- [34] In search of the neural circuits of intrinsic motivation View paper
- [35] Phasic dopamine as a prediction error of intrinsic and extrinsic reinforcements driving both action acquisition and reward maximization: A simulated robotic study View paper
- [36] Surprise signals in the supplementary eye field: rectified prediction errors drive exploration-exploitation transitions View paper
- [37] Predication-Error-Based Intrinsically Motivated Saccade Learning View paper
- [38] Contrastive difference predictive coding View paper
- [39] Curiosity-driven learning in artificial intelligence and its applications View paper
- [40] Multistep Quasimetric Learning for Scalable Goal-conditioned Reinforcement Learning View paper
- [41] Offline Goal-conditioned Reinforcement Learning with Quasimetric Representations View paper
- [42] Goal Reaching with Eikonal-Constrained Hierarchical Quasimetric Reinforcement Learning View paper
- [43] Quasimetric decision transformer: enhancing goal-conditioned reinforcement learning with structured distance guidance View paper