Novelty Assessment Report

Paper: Two-Way Is Better Than One: Bidirectional Alignment with Cycle Consistency for Exemplar-Free Class-Incremental Learning

PDF URL: https://openreview.net/pdf?id=7UfZAxKo5K

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Continual learning (CL) seeks models that acquire new skills without erasing prior knowledge. In exemplar-free class-incremental learning (EFCIL), this challenge is amplified because past data cannot be stored, making representation drift for old classes particularly harmful. Prototype-based EFCIL is attractive for its efficiency, yet prototypes drift as the embedding space evolves; thus, projection-based drift compensation has become a popular remedy. We show, however, that existing one-directional projections introduce systematic bias: they either retroactively distort the current feature geometry or align past classes only locally, leaving cycle inconsistencies that accumulate across tasks. We introduce bidirectional projector alignment during training: two maps, old$\to$new and new$\to$old, are trained during each new task with stop-gradient gating and a cycle-consistency objective so that transport and representation co-evolve. Analytically, we prove that the cycle loss contracts the singular spectrum toward unity in whitened space and that improved transport of class means/covariances yields smaller perturbations of classification log-odds, preserving old-class decisions and directly mitigating catastrophic forgetting. Empirically, across standard EFCIL benchmarks, our method achieves unprecedented reductions in forgetting while maintaining very high accuracy on new tasks, consistently outperforming state-of-the-art approaches.

Core Task Landscape

This paper addresses: **Exemplar-Free Class-Incremental Learning with Prototype Drift Compensation**

A total of **40 papers** were analyzed and organized into a taxonomy with **31 categories**.

Taxonomy Overview

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

- **Prototype Drift Compensation Mechanisms**
- **Prototype Calibration and Refinement Strategies**
- **Feature Representation Stabilization**
- **Pseudo-Replay and Synthetic Data Generation**
- **Specialized Architectural and Modular Approaches**
- **Heterogeneity-Aware and Distribution-Based Methods**
- **Prototype Evolution and Reminiscence Mechanisms**
- **Domain-Specific and Application-Oriented Extensions**
- **Interpretability and Semi-Supervised Extensions**
- **Old Class Reconstruction and Resurrecting Strategies**

Complete Taxonomy Tree

- Exemplar-Free Class-Incremental Learning with Prototype Drift Compensation Survey Taxonomy
- Prototype Drift Compensation Mechanisms
  - Bidirectional and Cycle-Consistent Alignment ★ (1 papers)
  - [0] Two-Way Is Better Than One: Bidirectional Alignment with Cycle Consistency for Exemplar-Free Class-Incremental Learning (Anon et al., 2026) View paper
  - Adaptive Prototype Correction and Reconstruction (4 papers)
  - [2] Exemplar-free class incremental learning for rotating machinery fault diagnosis via adaptive prototype correction and separation network (Zongzhen Ye, 2025) View paper
  - [4] Non-Exemplar Class-Incremental Learning via Prototype Correction and Hierarchical Regularization for Specific Emitter Identification (Dingzhao Li, 2025) View paper
  - [6] Non-exemplar class-incremental learning via adaptive old class reconstruction (Shao-Kun Wang, 2023) View paper
  - [39] Similarity-based prototype reconstruction and feature reorganization for non-exemplar class incremental learning. (Chao Zhou, n.d.) View paper
  - Learnable Drift Compensation (1 papers)
  - [5] Exemplar-free Continual Representation Learning via Learnable Drift Compensation (Gomez-Villa, 2024) View paper
  - Semantic Shift Estimation and Dual-Projection (2 papers)
  - [13] Semantic drift compensation for class-incremental learning (Lu Yu, 2020) View paper
  - [29] Semantic Shift Estimation via Dual-Projection and Classifier Reconstruction for Exemplar-Free Class-Incremental Learning (He Run, 2025) View paper
- Prototype Calibration and Refinement Strategies
  - Training-Free Prototype Calibration (1 papers)
  - [1] Few-shot class-incremental learning via training-free prototype calibration (Wang, 2023) View paper
  - Test-Time and Online Prototype Adaptation (2 papers)
  - [19] Restoring Forgotten Knowledge in Non-Exemplar Class Incremental Learning through Test-Time Semantic Evolution (Cao Xu-sheng, 2025) View paper

- [20] Non-exemplar Online Class-Incremental Continual Learning via Dual-Prototype Self-Augment and Refinement (Fushuo Huo, 2024) View paper
- Multi-View and Temporary Constraint Refinement (1 papers)
- [26] Multi-view prototype balance and temporary proxy constraint for exemplar-free class-incremental learning (Heng Tian, 2025) View paper

- Feature Representation Stabilization
  - Elastic Feature Consolidation (2 papers)
  - [22] EFC++: Elastic Feature Consolidation with Prototype Re-balancing for Cold Start Exemplar-free Incremental Learning (Simone Magistri, 2025) View paper
  - [35] Elastic Feature Consolidation for Cold Start Exemplar-free Incremental Learning (Simone Magistri, 2024) View paper
  - Prospective and Forward-Compatible Representation Learning (1 papers)
  - [7] Prospective representation learning for non-exemplar class-incremental learning (Wuxuan Shi, 2024) View paper
  - Feature Calibration and Separation (1 papers)
  - [14] Fcs: Feature calibration and separation for non-exemplar class incremental learning (Qiwei Li, 2024) View paper
- Pseudo-Replay and Synthetic Data Generation
  - Adversarial and Model Inversion-Based Replay (2 papers)
  - [18] Dual-consistency model inversion for non-exemplar class incremental learning (Zihuan Qiu, 2024) View paper
  - [25] Adversarial Pseudo-replay for Exemplar-free Class-incremental Learning (Honda, 2025) View paper
  - Diffusion-Based Feature Replay (1 papers)
  - [34] Diffusion Model Meets Non-Exemplar Class-Incremental Learning and Beyond (Zhang Jichuan, 2024) View paper
  - Retrospective Feature Synthesis (1 papers)
  - [15] Efficient Non-Exemplar Class-Incremental Learning with Retrospective Feature Synthesis (Bai Liang, 2024) View paper
  - Condensed Prototype Replay with Augmentation (1 papers)
  - [36] Condensed Prototype Replay for Class Incremental Learning (Kong, 2023) View paper
- Specialized Architectural and Modular Approaches
  - Vision Transformer Adaptation with Gated Mechanisms (2 papers)
  - [28] Exemplar-free Continual Learning of Vision Transformers via Gated Class-Attention and Cascaded Feature Drift Compensation (Marco Cotogni, 2022) View paper
  - Adapter Merging and Task Consolidation (1 papers)
  - [31] Adapter Merging with Centroid Prototype Mapping for Scalable Class-Incremental Learning (Takuma Fukuda, 2025) View paper
  - Prototype Consolidation with Ensemble Mechanisms (1 papers)
  - [40] ProCEED: Prototype Consolidation and Ensemble-based Exemplar-Free Deep Incremental Learning (D KANDEL, n.d.) View paper
- Heterogeneity-Aware and Distribution-Based Methods
  - Class Distribution Heterogeneity Exploitation (1 papers)
  - [11] Fecam: Exploiting the heterogeneity of class distributions in exemplar-free continual learning (Goswami, 2023) View paper
  - Distribution-Aware Knowledge Prototyping (1 papers)
  - [10] Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification (Kunlun Xu, 2024) View paper
  - Contrastive Prototype Learning with Drift Estimation (1 papers)
  - [17] Clear: Contrastive-prototype learning with drift estimation for resource constrained stream mining (Zhuoyi Wang, 2021) View paper
- Prototype Evolution and Reminiscence Mechanisms
  - Continual Prototype Evolution in Shared Latent Space (1 papers)
  - [38] Continual Prototype Evolution: Learning Online from Non-Stationary Data Streams (Matthias De Lange, 2021) View paper
  - Progressive Prototype Evolution with Dual-Forgetting Mitigation (1 papers)
  - [32] Progressive Prototype Evolving for Dual-Forgetting Mitigation in Non-Exemplar Online Continual Learning (Qiwei Li, 2024) View paper
  - Prototype Reminiscence and Asymmetric Knowledge Aggregation (1 papers)
  - [33] Prototype Reminiscence and Augmented Asymmetric Knowledge Aggregation for Non-Exemplar Class-Incremental Learning (Wuxuan Shi, 2023) View paper
  - Evanescent Representation Revival (1 papers)
  - [21] Bring evanescent representations to life in lifelong class incremental learning (Marco Toldo, 2022) View paper
- Domain-Specific and Application-Oriented Extensions
  - Graph-Based Incremental Learning (2 papers)
  - [9] Inductive graph few-shot class incremental learning (Yayong Li, 2025) View paper
  - [30] Instance-Prototype Affinity Learning for Non-Exemplar Continual Graph Learning (Song Lei, 2025) View paper
  - Video and Spatial-Temporal Incremental Learning (1 papers)
  - [24] CSTA: Spatial-Temporal Causal Adaptive Learning for Exemplar-Free Video Class-Incremental Learning (Huabin Liu, 2025) View paper
  - Domain-Specific Applications (1 papers)
  - [3] Exemplar-Free Class Incremental Learning for Traffic Sign Classification (Mei-Chen Liu, 2025) View paper
  - Cross-Subject and Cross-Domain Generalization (2 papers)
  - [16] Multivariate prototype representation for domain-generalized incremental learning (Peng Can, 2024) View paper
  - [23] Prototype-Guided Non-Exemplar Continual Learning for Cross-subject EEG Decoding (Dan Li, 2025) View paper
- Interpretability and Semi-Supervised Extensions
  - Interpretable Prototypical Part-Based Learning (1 papers)
  - [8] Icicle: Interpretable class incremental continual learning (Dawid Rymarczyk, 2023) View paper
  - Semi-Supervised Class-Incremental Learning (1 papers)
  - [27] Towards Non-Exemplar Semi-Supervised Class-Incremental Learning (Liu Wenzhuo, 2024) View paper
- Old Class Reconstruction and Resurrecting Strategies (1 papers)
  - [12] Resurrecting old classes with new data for exemplar-free continual learning (Dipam Goswami, 2024) View paper

Narrative

Core task: exemplar-free class-incremental learning with prototype drift compensation. The field addresses the challenge of learning new classes sequentially without storing past exemplars, while mitigating the drift that occurs when class prototypes shift as the feature extractor adapts to new data. The taxonomy reveals a rich landscape organized around ten major branches. Prototype Drift Compensation Mechanisms focus on alignment strategies that correct or stabilize prototypes across incremental steps, including bidirectional and cycle-consistent methods like Bidirectional Alignment Cycle[0]. Prototype Calibration and Refinement Strategies emphasize post-hoc adjustments using techniques such as Training-Free Prototype Calibration[1] and Adaptive Prototype Correction[2]. Feature Representation Stabilization aims to preserve old-class discriminability through regularization and consolidation, while Pseudo-Replay and Synthetic Data Generation leverage generative models to reconstruct old-class information. Additional branches cover Specialized Architectural and Modular Approaches, Heterogeneity-Aware and Distribution-Based Methods like Distribution-Aware Knowledge Prototyping[10], Prototype Evolution and Reminiscence Mechanisms such as Prototype Reminiscence[33], Domain-Specific extensions including Traffic Sign Incremental[3], and strategies for Old Class Reconstruction like Resurrecting Old Classes[12].

A central tension across these branches involves balancing stability and plasticity: methods must prevent catastrophic forgetting of old classes while accommodating new ones. Within Prototype Drift Compensation Mechanisms, Bidirectional Alignment Cycle[0] pursues cycle-consistent alignment to ensure forward and backward prototype consistency, contrasting with simpler unidirectional corrections found in works like Learnable Drift Compensation[5] or Semantic Drift Compensation[13]. This bidirectional approach sits at the intersection of alignment-based drift correction and feature-space consistency enforcement, sharing conceptual ground with calibration strategies such as Adaptive Prototype Correction[2] but emphasizing mutual alignment rather than one-way adjustment. Meanwhile, methods in Pseudo-Replay branches like Retrospective Feature Synthesis[15] and Dual-Consistency Model Inversion[18] tackle drift indirectly by regenerating old-class features, offering a complementary perspective. The original work's focus on cycle consistency positions it as a principled solution to drift, addressing a core challenge that resonates across multiple branches while offering a distinct geometric perspective on prototype alignment.

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 bidirectional alignment with cycle consistency constraints to maintain coherence between old and new feature spaces. Sibling subtopics address prototype drift through different mechanisms: adaptive correction using density/topology adjustments, learnable compensation modules, and semantic shift estimation via dual projections. All approaches aim to mitigate prototype drift in exemplar-free class-incremental learning but differ in their alignment strategies and correction mechanisms.

Similarities: - All subtopics address the fundamental problem of prototype drift in exemplar-free class-incremental learning - Multiple approaches employ projection-based mechanisms to align or correct feature spaces across learning phases - Semantic Shift Estimation and the original leaf both utilize dual-projection frameworks, though with different objectives (shift estimation vs. cycle consistency) - All methods operate without storing exemplars from previous tasks

Differences: - The original leaf enforces bidirectional cycle consistency constraints, while Learnable Drift Compensation uses explicit learnable modules and Adaptive Prototype Correction employs dynamic adjustment mechanisms - Adaptive Prototype Correction focuses on prototype-level adjustments (density, topology, reconstruction), whereas the original leaf operates on feature space alignment - Semantic Shift Estimation explicitly quantifies drift magnitude, while the original leaf implicitly addresses drift through consistency constraints - Learnable Drift Compensation introduces trainable parameters specifically for drift correction, contrasting with the original leaf's constraint-based approach

Suggested Search Directions: - Investigate hybrid approaches combining cycle consistency with learnable drift compensation modules - Explore methods that integrate semantic shift estimation with bidirectional alignment constraints - Examine whether adaptive prototype correction mechanisms can be incorporated into cycle-consistent frameworks

Sibling Subtopics

- **Adaptive Prototype Correction and Reconstruction** (leaves: 1, papers: 4)
- Scope: Methods that dynamically adjust prototypes using density-based reinforcement, topology correction, or reconstruction mechanisms.
- Exclude: Static prototype storage or feature-level regularization methods belong to other subcategories.
- **Learnable Drift Compensation** (leaves: 1, papers: 1)
- Scope: Approaches that learn explicit drift compensation modules or parameters to correct prototype displacement.
- Exclude: Methods using fixed projection rules or calibration without learnable components belong to other subcategories.
- **Semantic Shift Estimation and Dual-Projection** (leaves: 1, papers: 2)
- Scope: Methods estimating semantic drift magnitude through dual-projection frameworks or classifier reconstruction.
- Exclude: Single-projection methods or those without explicit shift estimation belong to other subcategories.

Contributions Analysis

Overall novelty summary. The paper introduces bidirectional projector alignment with cycle consistency to address prototype drift in exemplar-free class-incremental learning. It resides in the 'Bidirectional and Cycle-Consistent Alignment' leaf under 'Prototype Drift Compensation Mechanisms,' where it is currently the sole occupant among 40 papers across the taxonomy. This isolation suggests the specific combination of bidirectional projection and cycle consistency during training represents a relatively unexplored niche within the broader drift compensation landscape, which includes four other subcategories addressing prototype shift through alternative mechanisms.

The taxonomy reveals substantial activity in neighboring directions. 'Adaptive Prototype Correction and Reconstruction' contains four papers exploring density-based reinforcement and topology correction, while 'Semantic Shift Estimation and Dual-Projection' includes two works on dual-projection frameworks without cycle constraints. The 'Prototype Calibration and Refinement Strategies' branch offers complementary post-hoc adjustment methods, and 'Pseudo-Replay and Synthetic Data Generation' provides an alternative paradigm through feature regeneration. The paper's bidirectional approach bridges alignment-based drift correction with geometric consistency enforcement, occupying conceptual space between unidirectional projection methods and calibration strategies that lack mutual alignment guarantees.

Among 11 candidates examined, the bidirectional alignment contribution shows one refutable candidate from six examined, while the geometry-preserving transport mechanism found no refutations among five candidates. The theoretical analysis contribution was not tested against prior work. The limited search scope—11 total candidates rather than an exhaustive survey—means these statistics reflect top-K semantic matches and citation expansion, not comprehensive coverage. The bidirectional alignment's single refutation suggests some overlap exists within the examined subset, while the transport mechanism appears more distinctive among the candidates reviewed.

Based on the 11-candidate search, the work appears to occupy a sparse research direction with limited direct precedent in its specific technical approach. The taxonomy structure confirms that while drift compensation is a crowded area overall, the particular combination of bidirectional projection and cycle consistency during training has minimal representation. However, the analysis cannot rule out relevant work outside the examined candidate set, and the single refutation for the core contribution warrants careful examination of overlap boundaries.

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

Contribution 1: Bidirectional projector alignment with cycle consistency during training

Description: The authors propose learning two projection maps simultaneously during each task—one from old to new feature space and one from new to old—using stop-gradient operations and a cycle-consistency loss. This approach allows transport and representation to evolve together, addressing limitations of prior two-stage, one-directional methods.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Transfer learning via representation learning

URL: View paper

Brief Assessment

Transfer Representation Learning[41] mentions cycle-consistency in domain alignment contexts, but the provided context is too fragmentary to establish whether it addresses the same problem space (exemplar-free class-incremental learning with simultaneous old→new and new→old projectors trained with stop-gradient operations). The candidate appears focused on transfer learning and domain alignment rather than continual learning drift compensation.

2. Neuro-Symbolic Diffusion

URL: View paper

Brief Assessment

Neuro-Symbolic Diffusion[46] mentions 'cycle consistency loss' in passing but provides no context about bidirectional feature space alignment, projector learning, or continual learning frameworks. The retrieved text fragments are too sparse to establish any meaningful comparison with the original paper's contribution of learning two projection maps (old→new and new→old) simultaneously with stop-gradient operations during continual learning tasks.

3. Deepcollaboration: Collaborative generative and discriminative models for class incremental learning

URL: View paper

Prior Art Analysis

DeepCollaboration[42] demonstrates prior work on bidirectional training with cycle consistency in continual learning. The candidate paper explicitly proposes 'bi-directional training to enforce the cycle consistency' between feature and image domains, where 'the first cycle consistency enables a model to reconstruct raw images from the latent vectors' (g(e(x)) → x) and 'the other cycle consistency encourages the latent representation of generated images to be similar with the feature vectors' (e(g(f)) → f). This bidirectional cycle-consistent training is performed during the main training phase, not as a post-hoc step, directly addressing the same problem of feature space alignment across tasks that the original paper claims as novel.

Evidence

Evidence 1 - **Rationale**: Both papers propose bidirectional training with cycle consistency during the main training phase. The candidate explicitly describes two-way mappings (encoder E and generator G) trained jointly with cycle consistency objectives, similar to the original's old→new and new→old projectors. - **Original**: we introduce bidirectional projector alignment during training: two maps, old→new and new→old, are trained during each new task with stop-gradient gating and a cycle-consistency objective so that transport and representation co-evolve. - **Candidate**: to eliminate the domain mismatch between the generated images and the real images, inspired by the recent work on imageto-image translation (zhu et al. 2017a), we propose to use bi-directional training to enforce the cycle consistency. the first cycle consistency enables a model to reconstruct raw im...

Evidence 2 - **Rationale**: Both papers formulate bidirectional alignment losses that enforce consistency between forward and backward mappings. The candidate's objective enforces e(g(f)) → f, while the original enforces similar bidirectional constraints with stop-gradient operations. - **Original**: bidirectional alignment.we seek (i)backward compatibilityby makingz new projectable to the old space viad, and (ii) aforwardmapathat transports old prototypes into the current space used for evaluation-without draggingf t backward. concretely, lbi =‖d(z new)-z old‖2 2 +‖a(z old)-stopgrad(z new)‖2 2. - **Candidate**: the other cycle consistency encourages the latent representation of generated images to be similar with the feature vectors sampled from: e(g(f)) !f. we try to optimize the following objective: gde= argmin g;e max d lgan(g;d) +disldis(g;e).

Evidence 3 - **Rationale**: Both papers emphasize joint training of bidirectional mappings during the main optimization phase rather than as a post-hoc step. The candidate explicitly states that E and G are 'connected through a unified feature distribution design and bidirectional joint training' to solve distribution mismatch. - **Original**: our idea: from two-stage tonearsingle-stage transport.motivated by the limitations of twostage drift compensation, we propose bidirectional cycle consistency that evolves adapter training intothe main task optimization so that transport and representation co-evolve. - **Candidate**: the whole system consists of three parts, the discriminative model e, the generative model g, and the domain alignment module fd;ar;agg. the discriminative model e is in fact an encoder which can perform representation learning and create good feature embeddings for classification. the generative mod...

Evidence 4 - **Rationale**: Both papers use cycle consistency to prevent degeneracies in bidirectional mappings. The candidate introduces feature adaptors with cycle-consistency constraints to ensure domain-invariant features can be translated bidirectionally, similar to the original's cycle loss enforcing a∘d ≈ i and d∘a ≈ i. - **Original**: cycle consistency.whilel bi aligns both directions, it does not by itself prevent degeneracies (e.g., rank loss in weakly correlated directions). we therefore add a cycle loss that nudges the compositions toward identity on the data support: lcyc =‖a(d(znew))-stopgrad(z new)‖2 2 +‖d(a(zold))-stopgra... - **Candidate**: domain-invariant features should be able to be translated from one domain to the other (li et al. 2019). to this end, we propose to train feature adaptors ar and ag separately which are applied to features of real and generated images, fr and fg, to translate the features from one domain to the othe...

4. MB2C: Multimodal Bidirectional Cycle Consistency for Learning Robust Visual Neural Representations

URL: View paper

Brief Assessment

Multimodal Bidirectional Cycle[44] applies bidirectional cycle consistency to align EEG brain activity with visual stimuli using dual-GAN for cross-modal translation, not to continual learning prototype drift compensation. The technical domains and problem settings are fundamentally different.

5. MedPEFT-CL: Dual-Phase Parameter-Efficient Continual Learning with Medical Semantic Adapter and Bidirectional Memory Consolidation

URL: View paper

Brief Assessment

MedPEFT-CL[43] focuses on medical vision-language segmentation with bidirectional Fisher-memory coordination for continual learning, not on general exemplar-free class-incremental learning with bidirectional feature space projectors and cycle consistency as proposed in the original paper.

6. Unsupervised Domain Adaptation for Mobile Semantic Segmentation based on Cycle Consistency and Feature Alignment

URL: View paper

Brief Assessment

Cycle Consistency Alignment[45] applies cycle consistency to domain adaptation for semantic segmentation between synthetic and real images, not to continual learning feature space alignment. The technical context (domain shift vs. representation drift across tasks) and application domain differ fundamentally.

Contribution 2: Theoretical analysis of cycle loss and decision stability

Description: The authors provide theoretical guarantees showing that minimizing the cycle loss contracts singular values toward one in whitened feature space, and that reduced alignment and cycle errors lead to tighter bounds on classification log-odds perturbations, thereby preserving decisions on old classes.

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: Geometry-preserving transport mechanism for drift mitigation

Description: The method maintains the geometric structure of old-class representations even as the feature extractor adapts to new tasks, which reduces the tendency to favor recent classes and improves retention of previously learned knowledge.

This contribution was assessed against **5 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Federated Continual Learning with Bounded Forgetting via Diffusion-Based Generative Replay in Edge Computing

URL: View paper

Brief Assessment

Diffusion-Based Generative Replay[49] focuses on federated continual learning with generative replay in edge computing environments, not on geometry-preserving feature transport mechanisms for class-incremental learning. The minimal context provided does not demonstrate prior work on bidirectional cycle-consistent transport for drift compensation.

2. The Geometry of Abstraction: Continual Learning via Recursive Quotienting

URL: View paper

Brief Assessment

Geometry of Abstraction[48] addresses continual learning through recursive metric contraction on temporal manifolds, focusing on topological quotient maps and hierarchical depth. The original paper's geometry-preserving transport operates in feature space to mitigate recency bias in class-incremental learning, which is a distinct technical domain from manifold topology theory.

3. Generalization-Preserved Learning: Closing the Backdoor to Catastrophic Forgetting in Continual Deepfake Detection

URL: View paper

Brief Assessment

Generalization-Preserved Learning[47] addresses continual deepfake detection using hyperbolic visual alignment and gradient projection, not general class-incremental learning with prototype transport. The technical focus and application domain differ fundamentally from the original paper's geometry-preserving transport for class prototypes in EFCIL.

4. Exploring continual learning for robust memory efficient solutions and adversarial challenges: a dissertation in Engineering and Applied Science

URL: View paper

Brief Assessment

Continual Learning Dissertation[50] focuses on bias-robust methods and adversarial attacks in continual learning, not on geometry-preserving feature transport mechanisms for mitigating recency bias through bidirectional alignment.

5. Preserving Linear Separability in Continual Learning by Backward Feature Projection

URL: View paper

Brief Assessment

Backward Feature Projection[51] focuses on preserving linear separability through backward projection of new features to old feature space, while the original paper proposes bidirectional cycle-consistent transport between old and new spaces. The candidate's approach is unidirectional (new→old) rather than bidirectional with cycle consistency.

Appendix: Text Similarity Detection

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

References

- [0] Two-Way Is Better Than One: Bidirectional Alignment with Cycle Consistency for Exemplar-Free Class-Incremental Learning View paper
- [1] Few-shot class-incremental learning via training-free prototype calibration View paper
- [2] Exemplar-free class incremental learning for rotating machinery fault diagnosis via adaptive prototype correction and separation network View paper
- [3] Exemplar-Free Class Incremental Learning for Traffic Sign Classification View paper
- [4] Non-Exemplar Class-Incremental Learning via Prototype Correction and Hierarchical Regularization for Specific Emitter Identification View paper
- [5] Exemplar-free Continual Representation Learning via Learnable Drift Compensation View paper
- [6] Non-exemplar class-incremental learning via adaptive old class reconstruction View paper

- [7] Prospective representation learning for non-exemplar class-incremental learning View paper
- [8] Icicle: Interpretable class incremental continual learning View paper
- [9] Inductive graph few-shot class incremental learning View paper
- [10] Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification View paper
- [11] Fecam: Exploiting the heterogeneity of class distributions in exemplar-free continual learning View paper
- [12] Resurrecting old classes with new data for exemplar-free continual learning View paper
- [13] Semantic drift compensation for class-incremental learning View paper
- [14] Fcs: Feature calibration and separation for non-exemplar class incremental learning View paper
- [15] Efficient Non-Exemplar Class-Incremental Learning with Retrospective Feature Synthesis View paper
- [16] Multivariate prototype representation for domain-generalized incremental learning View paper
- [17] Clear: Contrastive-prototype learning with drift estimation for resource constrained stream mining View paper
- [18] Dual-consistency model inversion for non-exemplar class incremental learning View paper
- [19] Restoring Forgotten Knowledge in Non-Exemplar Class Incremental Learning through Test-Time Semantic Evolution View paper
- [20] Non-exemplar Online Class-Incremental Continual Learning via Dual-Prototype Self-Augment and Refinement View paper
- [21] Bring evanescent representations to life in lifelong class incremental learning View paper
- [22] EFC++: Elastic Feature Consolidation with Prototype Re-balancing for Cold Start Exemplar-free Incremental Learning View paper
- [23] Prototype-Guided Non-Exemplar Continual Learning for Cross-subject EEG Decoding View paper
- [24] CSTA: Spatial-Temporal Causal Adaptive Learning for Exemplar-Free Video Class-Incremental Learning View paper
- [25] Adversarial Pseudo-replay for Exemplar-free Class-incremental Learning View paper
- [26] Multi-view prototype balance and temporary proxy constraint for exemplar-free class-incremental learning View paper
- [27] Towards Non-Exemplar Semi-Supervised Class-Incremental Learning View paper
- [28] Exemplar-free Continual Learning of Vision Transformers via Gated Class-Attention and Cascaded Feature Drift Compensation View paper
- [29] Semantic Shift Estimation via Dual-Projection and Classifier Reconstruction for Exemplar-Free Class-Incremental Learning View paper
- [30] Instance-Prototype Affinity Learning for Non-Exemplar Continual Graph Learning View paper
- [31] Adapter Merging with Centroid Prototype Mapping for Scalable Class-Incremental Learning View paper
- [32] Progressive Prototype Evolving for Dual-Forgetting Mitigation in Non-Exemplar Online Continual Learning View paper
- [33] Prototype Reminiscence and Augmented Asymmetric Knowledge Aggregation for Non-Exemplar Class-Incremental Learning View paper
- [34] Diffusion Model Meets Non-Exemplar Class-Incremental Learning and Beyond View paper
- [35] Elastic Feature Consolidation for Cold Start Exemplar-free Incremental Learning View paper
- [36] Condensed Prototype Replay for Class Incremental Learning View paper
- [37] Gated Class-Attention with Cascaded Feature Drift Compensation for Exemplar-free Continual Learning of Vision Transformers View paper
- [38] Continual Prototype Evolution: Learning Online from Non-Stationary Data Streams View paper
- [39] Similarity-based prototype reconstruction and feature reorganization for non-exemplar class incremental learning. View paper
- [40] ProCEED: Prototype Consolidation and Ensemble-based Exemplar-Free Deep Incremental Learning View paper
- [41] Transfer learning via representation learning View paper
- [42] Deepcollaboration: Collaborative generative and discriminative models for class incremental learning View paper
- [43] MedPEFT-CL: Dual-Phase Parameter-Efficient Continual Learning with Medical Semantic Adapter and Bidirectional Memory Consolidation View paper
- [44] MB2C: Multimodal Bidirectional Cycle Consistency for Learning Robust Visual Neural Representations View paper
- [45] Unsupervised Domain Adaptation for Mobile Semantic Segmentation based on Cycle Consistency and Feature Alignment View paper
- [46] Neuro-Symbolic Diffusion View paper
- [47] Generalization-Preserved Learning: Closing the Backdoor to Catastrophic Forgetting in Continual Deepfake Detection View paper
- [48] The Geometry of Abstraction: Continual Learning via Recursive Quotienting View paper
- [49] Federated Continual Learning with Bounded Forgetting via Diffusion-Based Generative Replay in Edge Computing View paper
- [50] Exploring continual learning for robust memory efficient solutions and adversarial challenges: a dissertation in Engineering and Applied Science View paper
- [51] Preserving Linear Separability in Continual Learning by Backward Feature Projection View paper