

Novelty Assessment Report

Paper: Asymptotic analysis of shallow and deep forgetting in replay with neural collapse

PDF URL: <https://openreview.net/pdf?id=IdW0d0mRnG>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

A persistent paradox in Continual Learning is that neural networks often retain linearly separable representations of past tasks even when their output predictions fail. We formalize this distinction as the gap between deep (feature-space) and shallow (classifier-level) forgetting. We demonstrate that experience replay affects these two levels asymmetrically: while even minimal buffers anchor feature geometry and prevent deep forgetting, mitigating shallow forgetting requires substantially larger buffers. To explain this, we extend the Neural Collapse framework to sequential training. We theoretically model deep forgetting as a geometric drift toward out-of-distribution subspaces, proving that replay guarantees asymptotic separability. In contrast, we show that shallow forgetting stems from an under-determined classifier optimization: the strong collapse of buffer data leads to rank-deficient covariances and inflated means, blinding the classifier to the true population boundaries. Our work unifies continual learning with OOD detection and challenges the reliance on large buffers, suggesting that explicitly correcting the statistical artifacts of Neural Collapse could unlock robust performance with minimal replay.

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: **Asymmetric Effects of Replay on Feature Space and Classifier Forgetting**

A total of **50 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

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

- **Replay Mechanisms and Buffer Management**
- **Generative Replay and Synthetic Data**
- **Prototype and Representation-Based Methods**
- **Classifier and Decision Boundary Adaptation**
- **Theoretical Foundations and Asymmetric Forgetting Analysis**
- **Domain-Specific Continual Learning Applications**
- **Surveys, Reviews, and Methodological Frameworks**

Complete Taxonomy Tree

- Asymmetric Effects of Replay on Feature Space and Classifier Forgetting Survey Taxonomy
- Replay Mechanisms and Buffer Management
 - Memory Selection and Sampling Strategies
 - Saliency and Importance-Based Selection (3 papers)
 - [24] Saliency-driven experience replay for continual learning (Giovanni Bellitto, 2024) [View paper](#)
 - [34] Selective experience replay for lifelong learning (David Isele, 2018) [View paper](#)
 - [43] Three-way experience replay for the prediction under concept drift (J Wang, 2025) [View paper](#)
 - Distribution-Aware Selection (2 papers)
 - [36] Prototype-guided memory replay for continual learning (Ho, 2023) [View paper](#)
 - [46] Contrast Metric Enhancement Based on Memory Extraction for Online Class-incremental Learning in Image Classification (Fu, 2025) [View paper](#)
 - Task-Specific and Adaptive Replay (3 papers)
 - [6] Adaptive memory replay for network intrusion detection: Tackling data drift and catastrophic forgetting (Nasreen Fathima A H, 2025) [View paper](#)
 - [29] Overcoming Forgetting in Fine-Grained Urban Flow Inference via Adaptive Knowledge Replay (Xu, 2023) [View paper](#)
 - [45] Meta-Replay with Adaptive Feature Fusion: A Scalable, Memory-Efficient Framework for Continual Learning (Arash Khajooinejad, 2025) [View paper](#)
 - Replay Integration and Training Dynamics
 - Balanced Replay and Loss Modulation (2 papers)
 - [2] An analysis of best-practice strategies for replay and rehearsal in continual learning (Alexander Krawczyk, 2024) [View paper](#)
 - [5] Balanced destruction-reconstruction dynamics for memory-replay class incremental learning (Yuhang Zhou, 2024) [View paper](#)
 - Dynamic Replay Scheduling (2 papers)
 - [15] Class-incremental experience replay for continual learning under concept drift (Korycki, 2021) [View paper](#)
 - [39] Dynamic Replay Training for Class-Incremental Learning (Yan Yang, 2024) [View paper](#)
- Generative Replay and Synthetic Data
 - Diffusion-Based Generative Replay (4 papers)
 - [3] Beyond Synthetic Replays: Turning Diffusion Features into Few-Shot Class-Incremental Learning Knowledge (Kim Junsu, 2025) [View paper](#)

- [9] Class-Incremental Learning using Diffusion Model for Distillation and Replay (Quentin Jodelet, 2023) [View paper](#)
- [10] Deep Generative Replay With Denoising Diffusion Probabilistic Models for Continual Learning in Audio Classification (Hyeon Ju Lee, 2024) [View paper](#)
- [49] Federated Continual Learning with Bounded Forgetting via Diffusion-Based Generative Replay in Edge Computing (Zaobo He, 2025) [View paper](#)
- GAN-Based Generative Replay (3 papers)
- [17] Model-free generative replay for lifelong reinforcement learning: Application to starcraft-2 (Daniels, 2022) [View paper](#)
- [18] Bioslam: A bioinspired lifelong memory system for general place recognition (Peng Yin, 2023) [View paper](#)
- [25] Memory replay gans: Learning to generate new categories without forgetting (Chenshen Wu, 2018) [View paper](#)
- Feature-Level and Latent Generative Replay (3 papers)
- [27] Always Be Dreaming: A New Approach for Data-Free Class-Incremental Learning (Smith James, 2021) [View paper](#)
- [37] DeepCollaboration: Collaborative Generative and Discriminative Models for Class Incremental Learning (Cui Bo, 2021) [View paper](#)
- [40] Environmental Sound Classification Based on Continual Learning (Yadong Sun, 2023) [View paper](#)
- Prototype and Representation-Based Methods
 - Prototype Distillation and Alignment (3 papers)
 - [11] HyperKD: Lifelong Hyperspectral Image Classification With Cross-Spectral Spatial Knowledge Distillation (Zhenlin Li, 2025) [View paper](#)
 - [42] Few-Shot Lifelong Hyperspectral Image Classification With Dual-Branch Contrast and Geometric Alignment (Zhenlin Li, 2025) [View paper](#)
 - [50] Exemplar-Free Lifelong Hyperspectral Classification with Spectral-Spatial Prototype Alignment (Zhenlin Li, 2025) [View paper](#)
 - Dynamic Prototype Replay (2 papers)
 - [13] Disentangled Prototype-Guided Dynamic Memory Replay for Continual Learning in Acoustic Signal Classification (Seok-hun Choi, 2024) [View paper](#)
 - [22] Few Shot Class Incremental Learning via Efficient Prototype Replay and Calibration (Wei Zhang, 2023) [View paper](#)
 - Representation Space Maintenance (3 papers)
 - [23] Distribution-Level Memory Recall for Continual Learning: Preserving Knowledge and Avoiding Confusion (Shaoxu Cheng, 2024) [View paper](#)
 - [30] Adaptive feature consolidation residual network for exemplar-free continuous diagnosis of rotating machinery with fault-type increments (Yan Zhang, 2024) [View paper](#)
 - [33] Representation Space Maintenance: Against Forgetting in Continual Learning (Rui Niu, 2024) [View paper](#)
- Classifier and Decision Boundary Adaptation
 - Bias Correction and Calibration (2 papers)
 - [16] Signal Modulation Recognition via Bias Adjustment-Based Class Incremental Learning (Bo Xu, 2024) [View paper](#)
 - [21] Retrospective adversarial replay for continual learning (L Kumari, 2022) [View paper](#)
 - Decision Boundary Adjustment (1 papers)
 - [38] Mitigating Catastrophic Forgetting in Task-Incremental Continual Learning with Adaptive Classification Criterion (Luo Yun, 2023) [View paper](#)
- Theoretical Foundations and Asymmetric Forgetting Analysis ★ (2 papers)
 - [0] Asymptotic analysis of shallow and deep forgetting in replay with neural collapse (Anon et al., 2026) [View paper](#)
 - [35] Brain-inspired feature exaggeration in generative replay for continual learning (Jack Millichamp, 2021) [View paper](#)
- Domain-Specific Continual Learning Applications
 - Vision and Spatial Data
 - Object Detection and Segmentation (4 papers)
 - [7] Revisiting Generative Replay for Class Incremental Object Detection (Shizhou Zhang, 2025) [View paper](#)
 - [14] Beyond Background Shift: Rethinking Instance Replay in Continual Semantic Segmentation (Yin Hongmei, 2025) [View paper](#)
 - [26] Augmented Box Replay: Overcoming Foreground Shift for Incremental Object Detection (Yuyang Liu, 2023) [View paper](#)
 - [32] NIFF: Alleviating Forgetting in Generalized Few-Shot Object Detection via Neural Instance Feature Forging (Karim Guirguis, 2023) [View paper](#)
 - Image Classification (2 papers)
 - [12] Cross-Domain Incremental Feature Learning for ALS Point Cloud Semantic Segmentation With Few Samples (Mofan Dai, 2024) [View paper](#)
 - [28] Replay with Stochastic Neural Transformation for Online Continual EEG Classification (Tiehang Duan, 2023) [View paper](#)
 - Audio and Signal Processing (1 papers)
 - [41] AFT: An Exemplar-Free Class Incremental Learning Method for Environmental Sound Classification (Chen Xin-yi, 2025) [View paper](#)
 - Text and Natural Language (2 papers)
 - [31] InfoCL: Alleviating Catastrophic Forgetting in Continual Text Classification from An Information Theoretic Perspective (Song, 2023) [View paper](#)
 - [44] GeRe: Towards Efficient Anti-Forgetting in Continual Learning of LLM via General Samples Replay (Zhang Yunan, 2025) [View paper](#)
 - Spatio-Temporal and Sequential Data (2 papers)
 - [8] A Unified Replay-Based Continuous Learning Framework for Spatio-Temporal Prediction on Streaming Data (Hao Miao, 2024) [View paper](#)
 - [48] Adaptive Passenger Flow Prediction in Metro Systems: An Incremental Learning Framework for Evolving Patterns (Fei Wu, 2025) [View paper](#)
- Surveys, Reviews, and Methodological Frameworks (5 papers)
 - [1] Online continual learning: A systematic literature review of approaches, challenges, and benchmarks (Rezaee, 2025) [View paper](#)
 - [4] Towards Reliable Forgetting: A Survey on Machine Unlearning Verification, Challenges, and Future Directions (L Xue, 2025) [View paper](#)
 - [19] Exploring the current state and future potential of generative artificial intelligence using a generative artificial intelligence (Pescap, 2024) [View paper](#)
 - [20] NI-SSCL: A Neuroplasticity-Inspired Method for Semi-Supervised Continual Learning (Guanglei Xie, 2025) [View paper](#)
 - [47] Overcoming Catastrophic Forgetting: Geometric Techniques in Incremental Machine Learning (Nokhwal, 2025) [View paper](#)

Narrative

Core task: asymmetric effects of replay on feature space and classifier forgetting. The field of continual learning addresses how models can acquire new knowledge without catastrophically forgetting previous tasks. The taxonomy reveals a rich structure organized around several complementary themes. Replay Mechanisms and Buffer Management explores how stored examples are selected and maintained, with works like Replay Rehearsal Best Practices[2] and Adaptive Memory Replay Intrusion[6] examining storage strategies. Generative Replay and Synthetic Data investigates model-based approaches that synthesize past data using generative models, including diffusion-based methods such as Diffusion Distillation Replay[9] and Federated Diffusion Replay[49]. Prototype and Representation-Based Methods focuses on compact summaries and feature-level retention, while Classifier and Decision Boundary Adaptation tackles how decision surfaces evolve across tasks. Theoretical Foundations and Asymmetric Forgetting Analysis provides the analytical lens for understanding differential forgetting rates, and Domain-Specific Applications demonstrate these principles in specialized contexts like audio, hyperspectral imaging, and spatiotemporal prediction.

A central tension runs through the literature between memory efficiency and retention fidelity: some approaches store raw examples while others rely on generative reconstruction or prototype compression. Recent work increasingly examines the interplay between representation learning and classifier adaptation, recognizing that forgetting manifests differently at these two levels. Shallow Deep Forgetting[0] sits squarely within the Theoretical Foundations branch, providing analytical insight into how replay asymmetrically mitigates forgetting in feature extractors versus classification heads. This contrasts with neighboring work like Brain Inspired Feature Exaggeration[35], which emphasizes biologically motivated mechanisms for enhancing feature discriminability. While many studies focus on empirical replay strategies or generative alternatives, Shallow Deep Forgetting[0] offers a principled decomposition of forgetting dynamics, helping to explain why replay's protective effects vary across network depth and informing more targeted mitigation strategies.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Brain-inspired feature exaggeration in generative replay for continual learning

Authors: Jack Millichamp, Chen Xi, Xi Chen | **Year/Venue:** 2021 • arXiv.org | **URL:** [View paper](#)

Abstract

The catastrophic forgetting of previously learnt classes is one of the main obstacles to the successful development of a reliable and accurate generative continual learning model. When learning new classes, the internal representation of previously learnt ones can often be overwritten, resulting in the model's "memory" of earlier classes being lost over time. Recent developments in neuroscience have uncovered a method through which the brain avoids its own form of memory interference. Applying a t...

Relationship Analysis

Both papers belong to the Theoretical Foundations and Asymmetric Forgetting Analysis category, examining mechanisms of forgetting in continual learning. While the original paper provides a theoretical analysis of asymmetric forgetting between feature space and classifier using Neural Collapse theory and replay buffers, the candidate paper proposes a brain-inspired practical method that exaggerates feature differences through reconstruction repulsion to prevent memory interference. The key difference is that the original paper offers mathematical characterization of the deep-shallow forgetting gap, whereas the candidate paper introduces a novel loss function to mitigate forgetting by enhancing feature distinctiveness.

Contributions Analysis

Overall novelty summary. The paper formalizes the distinction between deep (feature-space) and shallow (classifier-level) forgetting in continual learning, demonstrating that replay affects these levels asymmetrically. It resides in the 'Theoretical Foundations and Asymmetric Forgetting Analysis' leaf, which contains only two papers in the entire 50-paper taxonomy. This sparse population suggests the paper addresses a relatively underexplored theoretical niche within continual learning, focusing on mechanistic analysis rather than empirical method development.

The taxonomy reveals that most continual learning research concentrates on practical techniques: replay mechanisms (11 papers across three sub-branches), generative methods (10 papers), and prototype-based approaches (8 papers). The paper's theoretical positioning contrasts sharply with these empirical branches. Its closest conceptual neighbors include classifier adaptation methods (3 papers) and representation maintenance work (3 papers), which address forgetting at specific network levels but lack the unified geometric framework proposed here. The taxonomy's domain-specific branch (13 papers) further highlights the field's applied orientation.

Among 29 candidates examined across three contributions, none clearly refuted the paper's claims. The formalization of deep versus shallow forgetting examined 9 candidates with no refutations; the Neural Collapse extension examined 10 with none; and the mechanistic explanation of shallow forgetting examined 10 with none. This absence of overlapping prior work within the limited search scope suggests the specific framing—linking Neural Collapse theory to continual learning's asymmetric forgetting—represents a novel analytical angle, though the search scale precludes definitive conclusions about the broader literature.

Based on top-29 semantic matches, the work appears to occupy a distinct theoretical position. The taxonomy structure confirms that analytical studies of forgetting mechanisms remain sparse compared to method-oriented research. However, the limited search scope means potentially relevant work in adjacent fields (e.g., neural collapse literature outside continual learning, or OOD detection theory) may not be fully captured. The novelty assessment reflects what was examined, not an exhaustive field survey.

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

Contribution 1: Formalization of deep versus shallow forgetting and asymmetric replay effects

Description: The authors formalize the distinction between deep forgetting (loss of feature-space separability) and shallow forgetting (classifier-level degradation). They empirically demonstrate that replay buffers mitigate these two forms of forgetting at fundamentally different rates, with small buffers sufficient for deep forgetting but large buffers required for shallow forgetting.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Forget-free continual learning with winning subnetworks

URL: [View paper](#)

Brief Assessment

Forget Free Winning Subnetworks[54] focuses on subnetwork-based continual learning without replay buffers, not on analyzing deep versus shallow forgetting or replay buffer effects. The candidate addresses catastrophic forgetting through architectural isolation rather than studying forgetting mechanisms.

2. Rewiring neurons in non-stationary environments

URL: [View paper](#)

Brief Assessment

Rewiring Neurons Non Stationary[58] focuses on structural plasticity through neuron permutation in continual reinforcement learning environments, not on analyzing deep versus shallow forgetting or replay buffer effects in continual learning.

3. Depth aware hierarchical replay continual learning for knowledge based question answering

URL: [View paper](#)

Brief Assessment

Depth Aware Hierarchical Replay[51] focuses on continual learning for knowledge-based question answering tasks, not on formalizing deep versus shallow forgetting or analyzing asymmetric replay buffer effects in neural collapse frameworks.

4. Mind the Gap: Layerwise Proximal Replay for Stable Continual Learning

URL: [View paper](#)

Brief Assessment

Layerwise Proximal Replay[57] focuses on the stability gap at task transitions in continual learning, not on the distinction between deep (feature-space) and shallow (classifier-level) forgetting or their asymmetric response to replay buffer sizes.

5. Incremental learning algorithm for anomaly detection applied to computed tomography scans in nuclear industry

URL: [View paper](#)

Brief Assessment

Incremental Anomaly Detection CT[56] focuses on automated defect detection in nuclear power plant tool inspection using classification-based supervised anomaly detection on CT scans. It does not address deep versus shallow forgetting concepts or experience replay mechanisms in continual learning.

6. Orchestrate latent expertise: Advancing online continual learning with multi-level supervision and reverse self-distillation

URL: [View paper](#)

Brief Assessment

Orchestrate Latent Expertise[52] focuses on multi-level supervision and self-distillation for online continual learning, not on formalizing deep versus shallow forgetting or analyzing asymmetric replay buffer effects on these two forms of forgetting.

7. Data Efficient Continual Learning of Large Language Model

URL: [View paper](#)

Brief Assessment

Data Efficient LLM[59] focuses on continual learning in large language models using causal inference to address catastrophic forgetting and spurious correlations. It does not address the distinction between deep (feature-space) and shallow (classifier-level) forgetting, nor does it analyze replay buffer effects on these two forms of forgetting.

8. Pretrained language model in continual learning: A comparative study

URL: [View paper](#)

Brief Assessment

Pretrained Language Model Continual[53] focuses on continual learning in NLP using pretrained language models, examining layer-wise forgetting patterns. It does not address the formalization of deep versus shallow forgetting in the context of neural collapse theory or the asymmetric replay buffer effects that are central to the original paper's contribution.

9. Position: Continual Learning Benefits from An Evolving Population over An Unified Model

URL: [View paper](#)

Brief Assessment

Evolving Population Continual[55] focuses on population-based evolutionary approaches to continual learning through neural architecture search, not on analyzing deep versus shallow forgetting mechanisms or replay buffer effects in feature versus classifier spaces.

Contribution 2: Extension of Neural Collapse framework to continual learning

Description: The authors extend the Neural Collapse theoretical framework to continual learning settings, including task-incremental, class-incremental, and domain-incremental learning. They characterize the asymptotic geometry of features and classifier heads under sequential training and prove that replay guarantees asymptotic separability in feature space.

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. Memory-efficient continual learning with neural collapse contrastive

URL: [View paper](#)

Brief Assessment

Neural Collapse Contrastive[69] focuses on using NC-based fixed prototypes for contrastive learning in continual learning, not on extending NC theory to characterize asymptotic geometry under sequential training as the original paper does.

2. Neural collapse terminus: A unified solution for class incremental learning and its variants

URL: [View paper](#)

Brief Assessment

Neural Collapse Terminus[62] focuses on class incremental learning with a fixed ETF structure as a consistent target, not on characterizing asymptotic geometry under sequential training or proving replay guarantees for feature separability as in the original paper.

3. Learn by Reasoning: Analogical Weight Generation for Few-Shot Class-Incremental Learning

URL: [View paper](#)

Brief Assessment

Learn by Reasoning[66] focuses on few-shot class-incremental learning using analogical weight generation and mentions neural collapse theory only for semantic conversion between prototypes and weights. It does not extend the Neural Collapse framework to continual learning settings or prove theoretical results about asymptotic geometry under sequential training.

4. Learning equi-angular representations for online continual learning

URL: [View paper](#)

Brief Assessment

Equi Angular Representations[61] focuses on inducing neural collapse in online continual learning with single-epoch training constraints, while the original paper extends neural collapse theory to analyze asymptotic geometry under sequential multi-epoch training across task-incremental, class-incremental, and domain-incremental settings. The candidate addresses practical online learning challenges (bias problem, preparatory data), whereas the original provides theoretical characterization of feature-space forgetting and replay dynamics.

5. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning

URL: [View paper](#)

Brief Assessment

Neural Collapse Feature Classifier[60] applies neural collapse to few-shot class-incremental learning (FSCIL) by pre-assigning fixed ETF classifier prototypes, whereas the original paper extends neural collapse theory to analyze asymptotic geometry under sequential training with replay in task-incremental, class-incremental, and domain-incremental settings. These are distinct contributions addressing different aspects of continual learning.

6. Mitigating non-representative prototypes and representation bias in few-shot continual relation extraction

URL: [View paper](#)

Brief Assessment

Mitigating Non Representative Prototypes[63] focuses on few-shot continual relation extraction using the General Orthogonal Frame (GOF) structure for class prototypes, not on extending Neural Collapse theory to analyze continual learning dynamics or proving replay guarantees for feature separability.

7. Sequential-in-time training of nonlinear parametrizations for solving time-dependent partial differential equations

URL: [View paper](#)

Brief Assessment

Sequential Time Training PDEs[68] focuses on numerical methods for solving time-dependent PDEs using neural networks, not on continual learning or Neural Collapse theory in sequential task learning settings.

8. Compress to One Point: Neural Collapse for Pre-Trained Model-Based Class-Incremental Learning

URL: [View paper](#)

Brief Assessment

Compress to One Point[65] applies Neural Collapse theory to pre-trained model adaptation for class-incremental learning, focusing on constructing optimal classifier structures. The original paper extends Neural Collapse to analyze forgetting dynamics in sequential training across task-incremental, class-incremental, and domain-incremental settings with replay mechanisms.

9. Normalization and effective learning rates in reinforcement learning

URL: [View paper](#)

Brief Assessment

Normalization Effective Learning Rates[67] focuses on normalization layers and effective learning rates in reinforcement learning, not on extending Neural Collapse theory to continual learning settings or analyzing feature geometry under sequential training.

10. Learning optimal inter-class margin adaptively for few-shot class-incremental learning via neural collapse-based meta-learning

URL: [View paper](#)

Brief Assessment

Optimal Inter Class Margin[64] applies Neural Collapse theory to few-shot class-incremental learning (FSCIL) for computing optimal inter-class margins, not for analyzing asymptotic geometry or proving replay guarantees in sequential training as the original paper does.

Contribution 3: Mechanistic explanation of shallow forgetting via under-determined classifier optimization

Description: The authors demonstrate that shallow forgetting arises because Neural Collapse causes buffer data to collapse into a low-dimensional subspace, creating rank-deficient covariances and inflated means. This renders classifier optimization under-determined, causing decision boundaries to misalign with true population boundaries despite preserved feature separability.

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. Consistent Low-Rank Adaptation of Two-layer Neural Networks: A Nonparametric Statistics Approach

URL: [View paper](#)

Brief Assessment

Consistent Low Rank Adaptation[79] focuses on low-rank adaptation for transfer learning in two-layer neural networks using closed-form solutions, not on continual learning or the shallow/deep forgetting distinction in replay-based settings.

2. Storing encoded episodes as concepts for continual learning

URL: [View paper](#)

Brief Assessment

Storing Encoded Episodes[77] focuses on autoencoder-based memory compression and pseudo-rehearsal for continual learning, not on analyzing classifier optimization dynamics or neural collapse phenomena that cause shallow forgetting.

3. Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification

URL: [View paper](#)

Brief Assessment

Distribution Aware Knowledge Prototyping[71] addresses lifelong person re-identification using distribution-aware prototypes for knowledge transfer, not the mechanistic analysis of shallow forgetting through rank-deficient covariances in classifier optimization that the original paper investigates.

4. Overcoming Catastrophic Forgetting for Fine-Tuning Pre-trained GANs

URL: [View paper](#)

Brief Assessment

Overcoming Forgetting GANs[72] focuses on fine-tuning pre-trained GANs and low-rank approximation structures in generator layers, not on continual learning classifier optimization or neural collapse phenomena.

5. Demystifying Language Model Forgetting with Low-rank Example Associations

URL: [View paper](#)

Brief Assessment

Demystifying Language Model Forgetting[78] focuses on predicting which upstream examples will be forgotten during LLM fine-tuning using low-rank matrix factorization, not on the mechanistic causes of shallow vs. deep forgetting in continual learning with neural collapse.

6. Learning to Optimize Resource Allocation in Dynamic Wireless Environments: Embracing the New While Engaging the Old

URL: [View paper](#)

Brief Assessment

Resource Allocation Dynamic Wireless[75] addresses resource allocation in wireless networks using continual learning techniques. It does not discuss neural collapse, classifier optimization, or the mechanistic causes of shallow forgetting in replay-based continual learning.

7. An Empirical Analysis of Forgetting in Pre-trained Models with Incremental Low-Rank Updates

URL: [View paper](#)

Brief Assessment

Forgetting Incremental Low Rank[76] focuses on forgetting in low-rank adapter updates for pretrained models across fine-grained tasks, not on the mechanistic explanation of shallow forgetting through rank-deficient covariances in classifier optimization that the original paper provides.

8. On the Implicit Adversariality of Catastrophic Forgetting in Deep Continual Learning

URL: [View paper](#)

Brief Assessment

Implicit Adversariality Forgetting[73] focuses on adversarial alignment between new-task updates and old-task high-curvature directions in continual learning, not on rank-deficient covariances in classifier optimization or Neural Collapse mechanisms.

9. EEC: Learning to encode and regenerate images for continual learning

URL: [View paper](#)

Brief Assessment

EEC Encode Regenerate[74] focuses on autoencoder-based memory compression and image regeneration for continual learning, not on analyzing classifier optimization dynamics or neural collapse phenomena that cause shallow forgetting.

10. Preserving principal subspaces to reduce catastrophic forgetting in fine-tuning

URL: [View paper](#)

Brief Assessment

Preserving Principal Subspaces[70] addresses catastrophic forgetting through SVD-based weight matrix preservation during fine-tuning, focusing on preventing knowledge loss by constraining updates to lower-impact subspaces. This differs fundamentally from the original paper's mechanistic analysis of shallow forgetting arising from rank-deficient covariances in replay buffers.

Appendix: Text Similarity Detection

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

References

- [0] Asymptotic analysis of shallow and deep forgetting in replay with neural collapse [View paper](#)
- [1] Online continual learning: A systematic literature review of approaches, challenges, and benchmarks [View paper](#)
- [2] An analysis of best-practice strategies for replay and rehearsal in continual learning [View paper](#)
- [3] Beyond Synthetic Replays: Turning Diffusion Features into Few-Shot Class-Incremental Learning Knowledge [View paper](#)
- [4] Towards Reliable Forgetting: A Survey on Machine Unlearning Verification, Challenges, and Future Directions [View paper](#)
- [5] Balanced destruction-reconstruction dynamics for memory-replay class incremental learning [View paper](#)
- [6] Adaptive memory replay for network intrusion detection: Tackling data drift and catastrophic forgetting [View paper](#)
- [7] Revisiting Generative Replay for Class Incremental Object Detection [View paper](#)
- [8] A Unified Replay-Based Continuous Learning Framework for Spatio-Temporal Prediction on Streaming Data [View paper](#)
- [9] Class-Incremental Learning using Diffusion Model for Distillation and Replay [View paper](#)
- [10] Deep Generative Replay With Denoising Diffusion Probabilistic Models for Continual Learning in Audio Classification [View paper](#)
- [11] HyperKD: Lifelong Hyperspectral Image Classification With Cross-Spectral Spatial Knowledge Distillation [View paper](#)
- [12] Cross-Domain Incremental Feature Learning for ALS Point Cloud Semantic Segmentation With Few Samples [View paper](#)
- [13] Disentangled Prototype-Guided Dynamic Memory Replay for Continual Learning in Acoustic Signal Classification [View paper](#)
- [14] Beyond Background Shift: Rethinking Instance Replay in Continual Semantic Segmentation [View paper](#)
- [15] Class-incremental experience replay for continual learning under concept drift [View paper](#)
- [16] Signal Modulation Recognition via Bias Adjustment-Based Class Incremental Learning [View paper](#)
- [17] Model-free generative replay for lifelong reinforcement learning: Application to starcraft-2 [View paper](#)
- [18] Bioslam: A bioinspired lifelong memory system for general place recognition [View paper](#)
- [19] Exploring the current state and future potential of generative artificial intelligence using a generative artificial intelligence [View paper](#)
- [20] NI-SSCL: A Neuroplasticity-Inspired Method for Semi-Supervised Continual Learning [View paper](#)
- [21] Retrospective adversarial replay for continual learning [View paper](#)
- [22] Few Shot Class Incremental Learning via Efficient Prototype Replay and Calibration [View paper](#)

- [23] Distribution-Level Memory Recall for Continual Learning: Preserving Knowledge and Avoiding Confusion [View paper](#)
- [24] Saliency-driven experience replay for continual learning [View paper](#)
- [25] Memory replay gans: Learning to generate new categories without forgetting [View paper](#)
- [26] Augmented Box Replay: Overcoming Foreground Shift for Incremental Object Detection [View paper](#)
- [27] Always Be Dreaming: A New Approach for Data-Free Class-Incremental Learning [View paper](#)
- [28] Replay with Stochastic Neural Transformation for Online Continual EEG Classification [View paper](#)
- [29] Overcoming Forgetting in Fine-Grained Urban Flow Inference via Adaptive Knowledge Replay [View paper](#)
- [30] Adaptive feature consolidation residual network for exemplar-free continuous diagnosis of rotating machinery with fault-type increments [View paper](#)
- [31] InfoCL: Alleviating Catastrophic Forgetting in Continual Text Classification from An Information Theoretic Perspective [View paper](#)
- [32] NIFF: Alleviating Forgetting in Generalized Few-Shot Object Detection via Neural Instance Feature Forging [View paper](#)
- [33] Representation Space Maintenance: Against Forgetting in Continual Learning [View paper](#)
- [34] Selective experience replay for lifelong learning [View paper](#)
- [35] Brain-inspired feature exaggeration in generative replay for continual learning [View paper](#)
- [36] Prototype-guided memory replay for continual learning [View paper](#)
- [37] DeepCollaboration: Collaborative Generative and Discriminative Models for Class Incremental Learning [View paper](#)
- [38] Mitigating Catastrophic Forgetting in Task-Incremental Continual Learning with Adaptive Classification Criterion [View paper](#)
- [39] Dynamic Replay Training for Class-Incremental Learning [View paper](#)
- [40] Environmental Sound Classification Based on Continual Learning [View paper](#)
- [41] AFT: An Exemplar-Free Class Incremental Learning Method for Environmental Sound Classification [View paper](#)
- [42] Few-Shot Lifelong Hyperspectral Image Classification With Dual-Branch Contrast and Geometric Alignment [View paper](#)
- [43] Three-way experience replay for the prediction under concept drift [View paper](#)
- [44] GeRe: Towards Efficient Anti-Forgetting in Continual Learning of LLM via General Samples Replay [View paper](#)
- [45] Meta-Replay with Adaptive Feature Fusion: A Scalable, Memory-Efficient Framework for Continual Learning [View paper](#)
- [46] Contrast Metric Enhancement Based on Memory Extraction for Online Class-incremental Learning in Image Classification [View paper](#)
- [47] Overcoming Catastrophic Forgetting: Geometric Techniques in Incremental Machine Learning [View paper](#)
- [48] Adaptive Passenger Flow Prediction in Metro Systems: An Incremental Learning Framework for Evolving Patterns [View paper](#)
- [49] Federated Continual Learning with Bounded Forgetting via Diffusion-Based Generative Replay in Edge Computing [View paper](#)
- [50] Exemplar-Free Lifelong Hyperspectral Classification with Spectral-Spatial Prototype Alignment [View paper](#)
- [51] Depth aware hierarchical replay continual learning for knowledge based question answering [View paper](#)
- [52] Orchestrate latent expertise: Advancing online continual learning with multi-level supervision and reverse self-distillation [View paper](#)
- [53] Pretrained language model in continual learning: A comparative study [View paper](#)
- [54] Forget-free continual learning with winning subnetworks [View paper](#)
- [55] Position: Continual Learning Benefits from An Evolving Population over An Unified Model [View paper](#)
- [56] Incremental learning algorithm for anomaly detection applied to computed tomography scans in nuclear industry [View paper](#)
- [57] Mind the Gap: Layerwise Proximal Replay for Stable Continual Learning [View paper](#)
- [58] Rewiring neurons in non-stationary environments [View paper](#)
- [59] Data Efficient Continual Learning of Large Language Model [View paper](#)
- [60] Neural collapse inspired feature-classifier alignment for few-shot class incremental learning [View paper](#)
- [61] Learning equi-angular representations for online continual learning [View paper](#)
- [62] Neural collapse terminus: A unified solution for class incremental learning and its variants [View paper](#)
- [63] Mitigating non-representative prototypes and representation bias in few-shot continual relation extraction [View paper](#)
- [64] Learning optimal inter-class margin adaptively for few-shot class-incremental learning via neural collapse-based meta-learning [View paper](#)
- [65] Compress to One Point: Neural Collapse for Pre-Trained Model-Based Class-Incremental Learning [View paper](#)
- [66] Learn by Reasoning: Analogical Weight Generation for Few-Shot Class-Incremental Learning [View paper](#)
- [67] Normalization and effective learning rates in reinforcement learning [View paper](#)
- [68] Sequential-in-time training of nonlinear parametrizations for solving time-dependent partial differential equations [View paper](#)
- [69] Memory-efficient continual learning with neural collapse contrastive [View paper](#)
- [70] Preserving principal subspaces to reduce catastrophic forgetting in fine-tuning [View paper](#)
- [71] Distribution-aware knowledge prototyping for non-exemplar lifelong person re-identification [View paper](#)
- [72] Overcoming Catastrophic Forgetting for Fine-Tuning Pre-trained GANs [View paper](#)
- [73] On the Implicit Adversariality of Catastrophic Forgetting in Deep Continual Learning [View paper](#)
- [74] EEC: Learning to encode and regenerate images for continual learning [View paper](#)
- [75] Learning to Optimize Resource Allocation in Dynamic Wireless Environments: Embracing the New While Engaging the Old [View paper](#)
- [76] An Empirical Analysis of Forgetting in Pre-trained Models with Incremental Low-Rank Updates [View paper](#)
- [77] Storing encoded episodes as concepts for continual learning [View paper](#)
- [78] Demystifying Language Model Forgetting with Low-rank Example Associations [View paper](#)
- [79] Consistent Low-Rank Adaptation of Two-layer Neural Networks: A Nonparametric Statistics Approach [View paper](#)