# Novelty Assessment Report

**Paper**: PACE: Pretrained Audio Continual Learning
**PDF URL**: https://openreview.net/pdf?id=k5PgSlNc4E
**Venue**: ICLR 2026 Conference Submission
**Year**: 2026
**Report Generated**: 2026-01-01

## Abstract

Audio is a fundamental modality for analyzing speech, music, and environmental sounds. While pretrained audio models have significantly advanced audio understanding, they remain fragile in real-world scenarios where data distributions evolve over time. In this work, we present the first systematic benchmark for audio continual learning (CL) with pretrained models (PTMs) and provide a comprehensive analysis of its unique challenges. Unlike in the vision domain where parameter-efficient fine-tuning (PEFT) has proven effective for CL, directly applying such strategies to audio leads to poor performance. This is due to a fundamental property of audio backbones: they emphasize low-level spectral details rather than structured semantics, resulting in severe upstream–downstream misalignment. Through extensive empirical analysis, we identify a promising technical route based on analytic classifiers with first-session adaptation (FSA), but also uncover two major limitations: representation saturation in coarse-grained scenarios and representation shifts in fine-grained scenarios. To address these challenges, we propose **PACE**, an innovative method that improves FSA via a regularized analytic classifier and introduces multi-session adaptation through adaptive subspace-orthogonal PEFT for better semantic alignment. Additionally, we design spectrogram-based boundary-aware perturbations to mitigate representation overlap and improve stability. Experiments across six diverse audio CL benchmarks demonstrate that PACE substantially outperforms state-of-the-art baselines, representing a significant step toward robust and scalable audio CL with PTMs.

## Core Task Landscape

This paper addresses: **Audio Continual Learning with Pretrained Models**
A total of **50 papers** were analyzed and organized into a taxonomy with **25 categories**.

### Taxonomy Overview

The research landscape has been organized into the following main categories:
- **Continual Learning Methodologies and Optimization**
- **Foundation Models and Transfer Learning**
- **Application Domains and Empirical Studies**
- **Research Methodology and Design Frameworks**
- **Institutional and Policy Objectives**

### Complete Taxonomy Tree

- Audio Continual Learning with Pretrained Models Survey Taxonomy
- Continual Learning Methodologies and Optimization
  - Multi-Objective Optimization Frameworks
  - Evolutionary and Population-Based Multi-Objective Algorithms (3 papers)
    - [9] Indicator-based Multi-objective Evolutionary Algorithms (Jesús Guillermo Falcón-Cardona, 2020) View paper
    - [16] A Review on Constraint Handling Techniques for Population-based Algorithms: from single-objective to multi-objective optimization (Iman Rahimi, 2022) View paper
    - [17] Multiâobjective optimization of hybrid renewable energy system based on biomass and fuel cells (M. Samy, 2020) View paper
  - Multi-Objective System Modeling and Survey Studies (3 papers)
    - [45] A Survey of Multi-Objective Optimization in Wireless Sensor Networks: Metrics, Algorithms, and Open Problems (Fei, 2016) View paper
    - [46] A survey on modeling and optimizing multi-objective systems (Jin-Hee Cho, 2017) View paper
    - [47] Intelligent decision method for main control parameters of tunnel boring machine based on multi-objective optimization of excavation efficiency and cost (Bin Liu, 2021) View paper
  - Multi-Objective Optimization in Specific Engineering Domains (2 papers)
    - [7] RORA: Reinforcement learning based optimal distributed resource allocation strategies in vehicular cognitive radio networks for 6G (Mani Shekhar Gupta, 2025) View paper
    - [32] A two-stage multi-objective optimization method for envelope and energy generation systems of primary and secondary school teaching buildings in China (Yizhe Xu, 2021) View paper
  - Single-Task Learning and Objective Training (1 papers)
  - [10] Complement objective training (Chen, 2019) View paper
  - Continual Learning with Pretrained Models ★ (1 papers)
  - [0] PACE: Pretrained Audio Continual Learning (Anon et al., 2026) View paper
  - Feature Selection and Data Quality in Machine Learning (1 papers)
  - [29] Multi-Objective Feature Selection With Missing Data in Classification (Xue Yu, 2021) View paper
- Foundation Models and Transfer Learning
  - Foundation Models in Vision and Pathology (1 papers)

◦ [37] Monitoring the athlete training response: subjective self-reported measures trump commonly used objective measures: a systematic review (A. Saw, 2016) View paper
◦ [39] Social equity â□ primary objective of the Parliament of the Republic of Moldova (Ion GUCEAC, 2024) View paper

## Narrative

The field of audio continual learning with pretrained models addresses the challenge of adapting large-scale audio representations to sequential tasks without catastrophic forgetting. The taxonomy reveals five major branches: Continual Learning Methodologies and Optimization focuses on algorithmic strategies for mitigating forgetting and enabling incremental updates; Foundation Models and Transfer Learning examines how pretrained representations can be leveraged and fine-tuned across domains; Application Domains and Empirical Studies explores real-world deployments in speech, music, and environmental sound recognition; Research Methodology and Design Frameworks encompasses experimental protocols and evaluation metrics; and Institutional and Policy Objectives addresses broader organizational considerations. Works like Foundation Models Pathology[19] and Foundation Models Robotics[27] illustrate how pretrained architectures are being adapted beyond their original domains, while methodological contributions such as Complement Objective Training[10] and Indicator-based MOEA[9] provide optimization frameworks that balance multiple learning objectives.

Recent efforts reveal a tension between parameter efficiency and task performance, with many studies exploring how to selectively update pretrained weights while preserving prior knowledge. PACE[0] situates itself within the Continual Learning with Pretrained Models branch, emphasizing practical adaptation strategies that build on frozen or partially frozen representations. This contrasts with approaches like RORA[7], which may prioritize architectural innovations, or works such as LLM Software Development[14] that focus on deployment pipelines rather than core learning dynamics. The landscape also shows growing interest in multi-objective formulations, as seen in WSN Multi-Objective Survey[45] and Multi-Objective Systems Survey[46], reflecting the need to simultaneously optimize accuracy, memory footprint, and computational cost. Open questions remain around how to best align pretrained features with continually arriving data distributions and whether domain-specific inductive biases can be injected without full retraining.

## 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 continual learning scenarios where pretrained models are adapted sequentially across multiple tasks while mitigating catastrophic forgetting. The sibling subtopics address orthogonal concerns: one deals with data preprocessing and feature engineering for classification quality, while the other emphasizes optimization strategies for single-task training objectives. Neither sibling addresses the temporal, multi-task nature or memory retention challenges central to continual learning.

**Similarities:** - All three subtopics involve adapting or training models to improve performance on specific tasks - Each area may leverage pretrained representations as starting points, though with different downstream goals - All concern optimization or adaptation strategies within supervised learning contexts

**Differences:** - The original leaf explicitly handles sequential task arrival and catastrophic forgetting, while siblings focus on static single-task or data-centric scenarios - Feature Selection and Data Quality emphasizes input data preprocessing and reliability, whereas the original leaf addresses model parameter updates over time - Single-Task Learning optimizes for a fixed objective without task transitions, contrasting with the original leaf's multi-task temporal dynamics - The original leaf requires parameter-efficient or memory-based strategies to retain past knowledge, a concern absent in the sibling subtopics

**Suggested Search Directions:** - Investigate whether feature selection methods can be adapted for continual learning to identify task-invariant features - Explore how single-task optimization techniques (e.g., complement objectives) might be extended to multi-task continual settings - Examine the role of data quality and missing data handling in continual learning scenarios where task distributions shift over time

### Sibling Subtopics

• **Feature Selection and Data Quality in Machine Learning** (leaves: 1, papers: 1)
• Scope: Studies on feature selection methods, handling missing data, or improving data reliability in classification tasks.
• Exclude: Excludes end-to-end learning without explicit feature selection; see Single-Task Learning and Objective Training.
• **Single-Task Learning and Objective Training** (leaves: 1, papers: 1)
• Scope: Methods focusing on primary training objectives, complement objectives, or single-task performance optimization.
• Exclude: Excludes multi-objective or continual learning methods; see Multi-Objective Optimization Frameworks or Continual Learning with Pretrained Models.

## Contributions Analysis

**Overall novelty summary.** The paper introduces the first systematic benchmark for audio continual learning with pretrained models and proposes PACE, a method addressing representation saturation and shifts. It resides in the 'Continual Learning with Pretrained Models' leaf, which contains no sibling papers in the taxonomy. This isolation suggests the research direction is relatively sparse within the surveyed literature, indicating that audio-specific continual learning with pretrained models has received limited prior attention compared to broader continual learning methodologies or foundation model applications in vision and robotics.

The taxonomy places this work within 'Continual Learning Methodologies and Optimization,' adjacent to multi-objective optimization frameworks and single-task learning branches. Neighboring leaves include 'Foundation Models in Vision and Pathology' and 'Foundation Models in Robotics,' which explore pretrained model adaptation in other modalities. The scope note for the parent branch emphasizes sequential learning and adaptive training strategies, while excluding domain-specific applications without methodological contributions. This positioning highlights that the paper bridges methodological innovation (PACE) with domain-specific challenges (audio's low-level spectral emphasis), distinguishing it from purely algorithmic or purely applied studies.

Among 22 candidates examined, none clearly refute the three main contributions. The benchmark contribution examined 10 candidates with zero refutable matches, the PACE method examined 2 candidates with zero refutations, and the challenge identification examined 10 candidates with zero refutations. This limited search scope suggests that within the top-K semantic matches and citation expansions, no prior work explicitly addresses audio continual learning benchmarks or the specific upstream-downstream misalignment problem. The absence of refutable candidates across all contributions, combined with the sparse taxonomy leaf, indicates the work occupies a relatively unexplored niche.

Based on the limited literature search of 22 candidates, the paper appears to address a gap in audio-specific continual learning with pretrained models. However, the analysis does not cover exhaustive searches across all continual learning or audio processing venues, and the taxonomy's sparsity in this leaf may reflect search limitations rather than absolute novelty. The methodological contributions (PACE, first-session adaptation) and empirical findings (representation saturation, spectral misalignment) seem distinct within the examined scope, though broader surveys might reveal related work in adjacent audio or continual learning communities.

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

# Contribution 1: First systematic benchmark for audio continual learning with pretrained models

**Description**: The authors construct the first comprehensive benchmark specifically designed to evaluate continual learning methods on pretrained audio models. This benchmark includes six diverse audio datasets spanning coarse-grained and fine-grained tasks, and reveals fundamental challenges unique to the audio domain such as upstream-downstream misalignment and severe representation shifts.

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. Characterizing continual learning scenarios and strategies for audio analysis
**URL**: View paper

**Brief Assessment**

Audio Continual Learning[53] focuses on traditional continual learning methods (EWC, LwF, SI, GEM, etc.) without pretrained models, whereas the original paper specifically addresses continual learning WITH pretrained audio models and their unique challenges like upstream-downstream misalignment.

---

### 2. MetaCLBench: Meta Continual Learning Benchmark on Resource-Constrained Edge Devices
**URL**: View paper

**Brief Assessment**

MetaCLBench[56] focuses on meta-continual learning with limited labeled samples across image and audio modalities, not on systematic benchmarking of continual learning methods specifically for pretrained audio models. The candidate addresses meta-learning scenarios rather than the upstream-downstream misalignment challenges in pretrained audio models that the original paper investigates.

---

### 3. Audiobench: A universal benchmark for audio large language models
**URL**: View paper

**Brief Assessment**

Audiobench[51] focuses on evaluating audio large language models (AudioLLMs) across diverse tasks like speech understanding and audio captioning, not on continual learning scenarios with pretrained models. The two works address fundamentally different research problems.

---

### 4. CL-MASR: A Continual Learning Benchmark for Multilingual ASR
**URL**: View paper

**Brief Assessment**

CL-MASR[52] focuses specifically on multilingual ASR continual learning, while the original paper addresses general audio continual learning across diverse tasks (environmental sounds, speaker identification, music). These are distinct problem domains with different challenges.

---

### 5. LLMs Can Evolve Continually on Modality for X-Modal Reasoning
**URL**: View paper

**Brief Assessment**

X-Modal Reasoning[60] focuses on multimodal large language models with continual learning across different modalities (image, video, audio, depth, point cloud), not specifically on audio continual learning with pretrained audio models. The benchmark and challenges are fundamentally different.

---

### 6. Ddgr: Continual learning with deep diffusion-based generative replay
**URL**: View paper

**Brief Assessment**

DDGR[58] focuses on continual learning for image classification using diffusion-based generative replay, not audio domain benchmarks. The candidate addresses catastrophic forgetting in vision tasks, while the original work constructs audio-specific benchmarks.

---

### 7. CLASS: Continual learning approach for speech super-resolution
**URL**: View paper

**Brief Assessment**

CLASS[59] focuses on bandwidth extension using continual learning with self-supervised pretraining, not on constructing systematic benchmarks for evaluating continual learning methods across diverse audio datasets with pretrained models.

---

### 8. Less forgetting for better generalization: Exploring continual-learning fine-tuning methods for speech self-supervised representations
**URL**: View paper

**Brief Assessment**

Speech Self-Supervised Continual[54] focuses on continual learning during fine-tuning of speech encoders to reduce forgetting of pretraining tasks, not on constructing comprehensive benchmarks for audio continual learning across diverse datasets and task types.

---

### 9. Ucil: An unsupervised class incremental learning approach for sound event detection
**URL**: View paper

**Brief Assessment**

UCIL[57] focuses on unsupervised class incremental learning for sound event detection, not on systematic benchmarking of continual learning methods with pretrained audio models across diverse datasets.

---

### 10. Few-shot continual learning for audio classification
**URL**: View paper

**Brief Assessment**

Few-shot Audio Classification[55] addresses few-shot continual learning for audio classification, focusing on expanding vocabulary with minimal examples. It does not construct systematic benchmarks for evaluating continual learning methods on pretrained audio models across diverse datasets.

## Contribution 2: PACE method for pretrained audio continual learning

**Description**: The authors introduce PACE, a novel continual learning framework that addresses audio-specific challenges through three key components: improved first-session adaptation with layer-aware tuning, multi-session adaptation using adaptive subspace-orthogonal parameter-efficient fine-tuning, and boundary-aware perturbations to enhance representation stability and discriminability.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

### 1. Understanding Incremental Learning with Closed-form Solution to Gradient Flow on Overparamerterized Matrix Factorization

**URL**: View paper

**Brief Assessment**

Incremental Learning Gradient Flow[61] focuses on theoretical analysis of matrix factorization and incremental learning in gradient flow, not on audio continual learning or parameter-efficient fine-tuning methods for pretrained models.

### 2. Efficient Knowledge Transfer and Adaptation for Speech and Beyond

**URL**: View paper

**Brief Assessment**

Speech Knowledge Transfer[62] focuses on class-incremental spoken language understanding and parameter-efficient fine-tuning for audio classification, not on the specific PACE framework with its three components (layer-aware tuning, adaptive subspace-orthogonal PEFT, and boundary-aware perturbations) for pretrained audio continual learning.

## Contribution 3: Identification of fundamental audio continual learning challenges

**Description**: The authors systematically analyze audio continual learning and discover that unlike vision, audio models suffer from representation saturation during early adaptation on coarse-grained tasks and severe representation shifts on fine-grained tasks due to the mismatch between pretraining objectives focused on low-level spectral details and downstream semantic requirements.

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. Spectral Collapse Drives Loss of Plasticity in Deep Continual Learning

**URL**: View paper

**Brief Assessment**

Spectral Collapse Plasticity[64] focuses on Hessian spectral collapse in deep neural networks during continual learning, not on audio-specific representation saturation or shifts. The candidate addresses general deep learning plasticity loss through curvature analysis, while the original paper identifies domain-specific challenges in audio continual learning related to upstream-downstream mismatch between pretraining objectives and downstream semantic requirements.

### 2. Complementary Domain Adaptation and Generalization for Unsupervised Continual Domain Shift Learning

**URL**: View paper

**Brief Assessment**

The candidate paper (Complementary Domain Adaptation[67]) addresses unsupervised continual domain shift learning in general settings, focusing on domain adaptation and generalization strategies. It does not discuss audio-specific challenges, representation saturation, or the mismatch between pretraining objectives and downstream semantic requirements that are central to the original paper's contribution.

### 3. Continual test-time adaptation for end-to-end speech recognition on noisy speech

**URL**: View paper

**Brief Assessment**

Test-time Speech Adaptation[63] focuses on test-time adaptation for ASR under noisy conditions, not on analyzing fundamental representation saturation and shifts in audio continual learning with pretrained models across coarse/fine-grained tasks.

### 4. Unsupervised Continual Domain Shift Learning with Multi-Prototype Modeling

**URL**: View paper

**Brief Assessment**

Multi-Prototype Modeling[66] addresses unsupervised continual domain shift learning in visual domains, focusing on domain-invariant vs. domain-specific representations. It does not examine audio-specific challenges like representation saturation or spectral-semantic misalignment in audio continual learning.

### 5. Alleviating Representational Shift for Continual Fine-tuning

**URL**: View paper

**Brief Assessment**

Representational Shift Fine-tuning[65] focuses on intermediate layer representational shifts disrupting batch normalization in vision models, not on audio-specific challenges like representation saturation or upstream-downstream mismatch between spectral pretraining and semantic downstream tasks.

### 6. Online Distillation with Continual Learning for Cyclic Domain Shifts

**URL**: View paper

**Brief Assessment**

Online Distillation Cyclic[69] focuses on online distillation with cyclic domain shifts in video/autonomous driving contexts, not on analyzing representation saturation or shifts specific to audio continual learning with pretrained models.

### 7. A Continual Learning System with Self Domain Shift Adaptation for Fake News Detection

**URL**: View paper

**Brief Assessment**

Fake News Domain Shift[71] addresses fake news detection in a continual learning framework with domain shift adaptation for text data, not audio. The candidate focuses on non-stationarity in fake news streams using word2vec and LSTM, which is fundamentally different from the original paper's analysis of representation saturation and shifts in audio models with spectral features.

### 8. Online Continual Learning Under Domain Shift
**URL**: View paper

**Brief Assessment**

Online Continual Domain Shift[72] addresses domain shift between training and test distributions in continual learning, not the specific audio representation challenges (saturation and shifts due to spectral-semantic mismatch) identified in the original paper.

### 9. Continual Learning for Dynamic Children's Speech Adaptation in Whisper ASR
**URL**: View paper

**Brief Assessment**

Children Speech Whisper[68] focuses on adapting ASR models to children's speech using continual learning techniques like EWC, not on analyzing fundamental audio continual learning challenges such as representation saturation and shifts across coarse/fine-grained tasks.

### 10. Multi-Head Distillation for Continual Unsupervised Domain Adaptation in Semantic Segmentation
**URL**: View paper

**Brief Assessment**

Multi-Head Distillation[70] addresses continual unsupervised domain adaptation in semantic segmentation for visual data, not audio continual learning. The candidate focuses on catastrophic forgetting in multi-target visual domain adaptation scenarios, whereas the original contribution specifically identifies representation saturation and shifts unique to audio models due to spectral-semantic misalignment.

## Appendix: Text Similarity Detection

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

## References

- [0] PACE: Pretrained Audio Continual Learning View paper
- [1] The main objectives of the study View paper
- [2] The primary objectives of the show View paper
- [3] The heterogeneity of mesenchymal stem cells: an important issue to be addressed in cell therapy View paper
- [4] Pragmatism as a research paradigm and its implications for social work research View paper
- [5] Objectives, design and main findings until 2020 from the Rotterdam Study View paper
- [6] Text algorithms in economics View paper
- [7] RORA: Reinforcement learning based optimal distributed resource allocation strategies in vehicular cognitive radio networks for 6G View paper
- [8] A guide to research methodology: An overview of research problems, tasks and methods View paper
- [9] Indicator-based Multi-objective Evolutionary Algorithms View paper
- [10] Complement objective training View paper
- [11] A survey of public datasets for computer vision tasks in precision agriculture View paper
- [12] Estimated Research and Development Investment Needed to Bring a New Medicine to Market, 2009-2018. View paper
- [13] Research Objective and Methodology View paper
- [14] From Idea to Implementation: Evaluating the Influence of Large Language Models in Software Development - An Opinion Paper View paper
- [15] Designing a Web-Based Mail Management System at the Beringin Helvetia Sub-district Office View paper
- [16] A Review on Constraint Handling Techniques for Population-based Algorithms: from single-objective to multi-objective optimization View paper
- [17] Multiâobjective optimization of hybrid renewable energy system based on biomass and fuel cells View paper
- [18] A novel approach to traffic modelling based on road parameters, weather conditions and GPS data using feedforward neural networks View paper
- [19] Foundation Models -- A Panacea for Artificial Intelligence in Pathology? View paper
- [20] Redefining the primary objective of phase I oncology trials View paper
- [21] Recovery curve for patient reported outcomes and objective physical activity after primary total knee arthroplastyâa multicenter study using wearable technology View paper
- [22] The main objectives of the association View paper
- [23] Machine Learning-Based Drone Detection and Classification: State-of-the-Art in Research View paper
- [24] An objective definition for the main sequence of star-forming galaxies View paper
- [25] â dosing for pembrolizumab in the treatment of non-small cell lung cancer: A nationwide retrospective cohort study with a non-inferiority primary objective View paper
- [26] On the need to adjust the main objectives of the Bank of Russia View paper
- [27] Leveraging foundation models in robotics: Transforming task planning and contextual execution View paper
- [28] Euclid: Early Release Observations. Overview of the Perseus cluster and analysis of its luminosity and stellar mass functions View paper
- [29] Multi-Objective Feature Selection With Missing Data in Classification View paper
- [30] Assessing the effectiveness of entrepreneurship education in the universities of Tehran province based on an entrepreneurial intention model View paper
- [31] The distinctive features of a feasibility study: objectives and guiding questions View paper
- [32] A two-stage multi-objective optimization method for envelope and energy generation systems of primary and secondary school teaching buildings in China View paper
- [33] Addressing social problems through qualitative research View paper
- [34] Development, kinetic analysis, and economic feasibility of different Corn Stover-driven biorefineries to produce biohydrogen, bioethanol, and biomethane: A â View paper
- [35] Analysis of Export Competitiveness of Agri-Food Products at the EU-27 Level through the Perspective of Technical Complexity View paper
- [36] The main objective of the legal clinic: protecting the rights of the poor or skills of students View paper
- [37] Monitoring the athlete training response: subjective self-reported measures trump commonly used objective measures: a systematic review View paper

- [38] Technology Acceptance Model in M-learning context: A systematic review View paper
- [39] Social equity â primary objective of the Parliament of the Republic of Moldova View paper
- [40] Moroccan Medicinal plants as inhibitors against SARS-CoV-2 main protease: Computational investigations View paper
- [41] Research problems and techniques View paper
- [42] Unveiling the impact of psychological factors on consumer purchase intentions for overall sustainable success in green residential buildings: using SEM-ANN analysis View paper
- [43] Research Objectives and Method View paper
- [44] The Effect of Electronic Modules Based on Problem-Based Learning on Creative Thinking Ability View paper
- [45] A Survey of Multi-Objective Optimization in Wireless Sensor Networks: Metrics, Algorithms, and Open Problems View paper
- [46] A survey on modeling and optimizing multi-objective systems View paper
- [47] Intelligent decision method for main control parameters of tunnel boring machine based on multi-objective optimization of excavation efficiency and cost View paper
- [48] Combined effects of environmental factors on human perception and objective performance: A review of experimental laboratory works View paper
- [49] ChatGPT: ethical concerns and challenges in academics and research. View paper
- [50] The Rotterdam Study: objectives and design update View paper
- [51] Audiobench: A universal benchmark for audio large language models View paper
- [52] CL-MASR: A Continual Learning Benchmark for Multilingual ASR View paper
- [53] Characterizing continual learning scenarios and strategies for audio analysis View paper
- [54] Less forgetting for better generalization: Exploring continual-learning fine-tuning methods for speech self-supervised representations View paper
- [55] Few-shot continual learning for audio classification View paper
- [56] MetaCLBench: Meta Continual Learning Benchmark on Resource-Constrained Edge Devices View paper
- [57] Ucil: An unsupervised class incremental learning approach for sound event detection View paper
- [58] Ddgr: Continual learning with deep diffusion-based generative replay View paper
- [59] CLASS: Continual learning approach for speech super-resolution View paper
- [60] LLMs Can Evolve Continually on Modality for X-Modal Reasoning View paper
- [61] Understanding Incremental Learning with Closed-form Solution to Gradient Flow on Overparamerterized Matrix Factorization View paper
- [62] Efficient Knowledge Transfer and Adaptation for Speech and Beyond View paper
- [63] Continual test-time adaptation for end-to-end speech recognition on noisy speech View paper
- [64] Spectral Collapse Drives Loss of Plasticity in Deep Continual Learning View paper
- [65] Alleviating Representational Shift for Continual Fine-tuning View paper
- [66] Unsupervised Continual Domain Shift Learning with Multi-Prototype Modeling View paper
- [67] Complementary Domain Adaptation and Generalization for Unsupervised Continual Domain Shift Learning View paper
- [68] Continual Learning for Dynamic Children's Speech Adaptation in Whisper ASR View paper
- [69] Online Distillation with Continual Learning for Cyclic Domain Shifts View paper
- [70] Multi-Head Distillation for Continual Unsupervised Domain Adaptation in Semantic Segmentation View paper
- [71] A Continual Learning System with Self Domain Shift Adaptation for Fake News Detection View paper
- [72] Online Continual Learning Under Domain Shift View paper