

UDC 004.032.26:004.89
DOI [https://doi.org/10.15589/znp2026.1\(503\).2.9](https://doi.org/10.15589/znp2026.1(503).2.9)

DYNAMIC VALIDATION AND RELEVANCE MAINTENANCE OF NEURAL CORRELATES UNDER CONTINUAL LEARNING

ДИНАМІЧНА ВАЛІДАЦІЯ ТА ПІДТРИМКА РЕЛЕВАНТНОСТІ НЕЙРОННИХ КОРЕЛЯТИВ В УМОВАХ БЕЗПЕРЕРВНОГО НАВЧАННЯ

Oleksandr S. Verbytskyi
soomoalex@gmail.com
ORCID: 0009-0000-6848-8880

О. С. Вербицький,
магістр комп'ютерних наук

Admiral Makarov National University of Shipbuilding, Mykolaiv

Національний університет кораблебудування імені адмірала Макарова, м. Миколаїв

Abstract. Deep Neural Networks (DNNs) deployed in real-world, non-stationary environments frequently encounter performance degradation due to evolving data distributions, a phenomenon known as concept drift. Traditional adaptation strategies, such as full network retraining, are often computationally prohibitive and highly susceptible to the effect of catastrophic forgetting. There is an urgent scientific and engineering need for precise, interpretable, and resource-efficient mechanisms that can maintain model relevance in dynamic scenarios without the massive overhead of global parameter updates.

Objective. The primary goal of this research is to develop and evaluate a comprehensive framework for the dynamic validation and targeted maintenance of concept-specific neural correlates within deep learning architectures. This involves creating a system capable of autonomous knowledge decay detection at the level of individual functional units and performing surgical weight updates to restore predictive accuracy without disrupting the global knowledge integrity of the model.

Methodology. The proposed framework integrates L_0 -optimization for the initial topological localization of minimal neuron sets responsible for specific semantic concepts. To monitor the relevance of these localized correlates over time, a novel "Concept Adequacy Score" (A) is introduced, which measures the expected prediction confidence on incoming data streams using probabilistic methods. Upon detecting drift—formally defined as the score falling below a critical threshold τ —a localized adaptation mechanism is triggered. This mechanism employs targeted fine-tuning by isolating the parameters of identified neurons and freezing the remainder of the network's weights. The effectiveness of this approach was rigorously tested using a VGG16 architecture on sequentially shifting image distributions, simulating realistic drift scenarios.

Results. Experimental simulations demonstrate that the proposed framework successfully mitigates model degradation, maintaining classification accuracy at 93.1% in the final drift stage, compared to 96.5% for full retraining and only 51.2% for a static baseline. Crucially, the computational cost of the localized update was measured at approximately 0.8 GFLOPS per stage, representing a 94.8% reduction in operational overhead compared to the 15.5 GFLOPS required for full retraining. The results confirm that the system is resilient against catastrophic forgetting by preserving the stability of non-involved neurons while effectively integrating new conceptual variations.

Original contributions. This research establishes a novel intersection between Explainable AI (XAI) and Continual Learning (CL) by formalizing concept drift as a localized neural phenomenon rather than a global network failure. The key original contribution is the mathematical formulation of the Concept Adequacy Score and the subsequent development of a "surgical" adaptation protocol that treats neurons as independently updatable functional units. This methodology shifts the focus from global parameter optimization to localized architectural maintenance.

Practical significance. The developed framework is highly significant for AI applications in edge computing, autonomous robotics, and IoT systems where computational energy and memory are limited. By drastically reducing the resource costs associated with model maintenance, this approach enables the deployment of more sustainable, transparent, and adaptive intelligent systems. It provides a scalable solution for long-term AI operation in ever-changing environments, ensuring that models remain accurate throughout their entire lifecycle.

Key words: continual learning; concept drift; explainable AI (XAI); neural correlates; localized fine-tuning; resource-efficient AI.

Анотація. Глибокі нейронні мережі (DNN), що розгортаються в реальних нестационарних середовищах, часто стикаються з суттєвою деградацією продуктивності через зміну розподілу даних – явище, відоме як дрейф концепту. Традиційні стратегії адаптації, такі як повне перенавчання мережі на нових наборах даних, зазвичай є обчислювально витратними та вразливими до ефекту катастрофічного забування раніше набутих знань. На сьогодні існує нагальна наукова потреба у розробці точних, інтерпретованих та ресурсоефективних механізмів, здатних підтримувати релевантність моделі в динамічних сценаріях без необхідності глобального оновлення всіх параметрів системи.

Мета. Основною метою даного дослідження є розробка та експериментальна оцінка комплексного фреймворку для динамічної валідації та цільової підтримки концепт-специфічних нейронних корелятив у глибоких архітектурах. Це передбачає створення системи, здатної до автономного виявлення застарівання знань на рівні окремих функціональних одиниць мережі та виконання «хірургічних» оновлень ваг для відновлення точності прогнозування без порушення цілісності загальної бази знань моделі.

Методологія. Запропонований фреймворк поєднує методи L_0 -оптимізації для початкової топологічної локалізації мінімальних набірів нейронів, відповідальних за конкретні семантичні концепти. Для безперервного моніторингу релевантності цих локалізованих корелятив у часі введено нову метрику – «Оцінку адекватності концепту» (A), яка вимірює очікувану впевненість прогнозування на вхідних потоках даних. При виявленні дрейфу, що формально визначається падінням оцінки нижче критичного порогу τ , активується механізм локалізованої адаптації. Цей механізм використовує цільове тонке налаштування шляхом ізоляції параметрів ідентифікованих нейронів та повного заморожування решти ваг мережі. Ефективність підходу було протестовано на архітектурі VGG16 в умовах послідовної зміни розподілу зображень.

Результати. Експериментальне моделювання продемонструвало, що запропонований фреймворк успішно запобігає деградації моделі, підтримуючи точність класифікації на рівні 93,1% на фінальній стадії дрейфу, порівняно з 96,5% при повному перенавчанні та лише 51,2% для базової статичної моделі. Важливо, що обчислювальні витрати на локалізоване оновлення склали приблизно 0,8 GFLOPS на етап, що означає скорочення операційних витрат на 94,8% порівняно з 15,5 GFLOPS, необхідними для повного перенавчання. Результати підтверджують стійкість системи до катастрофічного забування завдяки стабільності не задіяних в оновленні нейронів.

Наукова новизна. Це дослідження встановлює новий теоретичний зв'язок між сферами пояснюваного ШІ (XAI) та безперервного навчання (Continual Learning), формалізуючи дрейф концепту як локалізований нейронний феномен, а не глобальну помилку всієї мережі. Ключовим внеском є математичне обґрунтування Оцінки адекватності концепту та розробка протоколу хірургічної адаптації, який розглядає нейрони як незалежно оновлювані функціональні одиниці. Такий підхід змінює парадигму з глобальної оптимізації параметрів на локалізоване архітектурне обслуговування.

Практичне значення. Розроблений фреймворк має критичне значення для застосувань ШІ в межах обчислень (edge computing), автономній робототехніці та системах інтернету речей, де ресурси пам'яті та енергії обмежені. Завдяки значному зниженню обчислювальних витрат на підтримку моделей, цей підхід дозволяє розгортати більш екологічні, прозорі та адаптивні інтелектуальні системи. Він пропонує масштабоване рішення для довгострокової експлуатації ШІ в реальних умовах, що постійно змінюються, гарантуючи точність моделей протягом усього їхнього життєвого циклу.

Ключові слова: безперервне навчання; дрейф концепту; пояснюваний ШІ (XAI); нейронні кореляти; локалізоване тонке налаштування; ресурсоефективний ШІ.

PROBLEM STATEMENT

The deployment of Deep Neural Networks (DNNs) in real-world, non-stationary environments exposes a fundamental flaw in traditional machine learning paradigms: the assumption of static data distributions. When a model trained on a specific dataset D_0 is exposed to a continuous stream of new data D_t where the statistical properties of the target variable or input features change, the model experiences "concept drift" [1]. In such scenarios, the complex correlations learned by the internal representations of the DNN rapidly become outdated, leading to severe model degradation.

The traditional approach to mitigating concept drift is full network retraining—combining historical data

with new data to update all network parameters Θ . This approach is fundamentally flawed for two reasons. First, it is computationally prohibitive, requiring massive energy and time expenditures that are incompatible with real-time or resource-constrained edge systems [2]. Second, if the model is retrained only on the new data stream D_t , it suffers from "catastrophic forgetting," a phenomenon where the network aggressively overwrites previously learned representations, destroying its historical knowledge base. Therefore, the critical problem lies in finding a mechanism to maintain model relevance and adapt to concept drift dynamically, without incurring the computational penalty of full retraining and without triggering catastrophic forgetting [3].

RELATED WORK

The challenge of model adaptation spans several sub-fields of artificial intelligence.

Continual Learning (CL). Traditional CL methods attempt to solve catastrophic forgetting through three main avenues: architectural allocation (dynamically adding branches to the network), regularization-based methods (e.g., Elastic Weight Consolidation, which penalizes changes to weights deemed important for old tasks), and replay-based methods (storing a buffer of old data to mix with new data). While effective, these methods often treat the network as a global "black box" and incur significant memory or computational overhead [4].

Explainable AI (XAI) and Network Sparsity. Recent advances in XAI have moved beyond simple saliency maps to understanding individual neural correlates [5]. Techniques utilizing L_0 -regularization allow for the precise localization of minimal sub-networks responsible for specific semantic concepts (e.g., isolating the specific filters in a Convolutional Neural Network that activate exclusively for "fur" or "wheels").

Parameter-Efficient Fine-Tuning (PEFT). Methods like LoRA (Low-Rank Adaptation) or Adapters freeze the pre-trained model and inject trainable layers to reduce update costs. However, PEFT methods are generally task-agnostic [6] and do not explicitly leverage the internal semantic structure (concept neurons) of the model for targeted updates.

Identified Gaps. Despite significant advancements, a crucial gap remains at the intersection of Continual Learning and Explainable AI. Existing CL methodologies update the network globally or semi-globally, leading to inefficiencies because they do not isolate *what* specifically has been forgotten or corrupted. Conversely, XAI techniques successfully identify concept-specific neural correlates but treat them as static endpoints for human interpretation, failing to utilize them as dynamic, updatable functional units.

There is currently no unified framework that leverages the interpretable, localized structure of a network to perform targeted, isolated updates—specifically, a system that detects decay in a localized neural concept and updates *only* those specific neurons, leaving the rest of the network untouched [7].

RESEARCH OBJECTIVES

To address the identified gaps, this research aims to develop a framework for the dynamic validation and targeted maintenance of concept-specific neural correlates under continual learning conditions. The specific objectives are:

1. Formulate a quantitative validation metric capable of continuously evaluating the relevance of localized neural correlates against evolving data streams in real-time.

2. Develop a localized adaptation algorithm that, upon detecting concept drift, fine-tunes only the outdated

conceptual neurons, explicitly bypassing the need for full global backpropagation.

3. Empirically validate the framework by simulating continuous concept drift, demonstrating that the proposed method maintains classification accuracy while reducing computational costs (measured in GFLOPS) by at least an order of magnitude compared to full retraining baselines.

Methodology, Study Object and Subject.

Object of study: The process of knowledge retention and adaptation in deep neural networks operating under continuous concept drift.

Subject of study: Mathematical models, algorithms, and metrics for the dynamic validation and localized parameter adaptation of concept-specific neural correlates.

Methodology: The research methodology integrates mathematical optimization (L_0 -regularization) for the initial topological localization of neurons, probabilistic and statistical methods for the formulation of the Concept Adequacy Score, and gradient-based targeted optimization for model fine-tuning. Empirical validation is conducted using controlled simulations of sequential data drift on standard image classification architectures (VGG16).

TECHNICAL APPROACH

The proposed technical framework operates as a continuous, autonomous cycle of localization, validation, and targeted adaptation. It shifts the paradigm of network updates from a global parameter space to a localized, concept-specific subspace [8].

1. Initial Concept Localization via L_0 -Optimization. Before dynamic maintenance can occur, the system must map the neural correlates of the base model. Given a model M with parameter space Θ trained on an initial dataset D_0 , we identify a minimal subset of neurons S_c that are highly responsive to a specific concept c (e.g., "Persian cat features"). We achieve this by applying L_0 -regularization during a brief localization phase, which forces the network to prune irrelevant activations, leaving only the dense core of neurons mathematically responsible for the concept.

2. Dynamic Validation and Drift Detection. Once deployed in a dynamic environment receiving data stream D_t , the model must self-diagnose. We introduce the **Concept Adequacy Score (A)**, a probabilistic metric that measures the expected confidence of the localized neuron set S_c in recognizing its target concept within the new data:

$$A(S_c, D_t) = \mathbb{E}_{x \in D_t} [P(c | M(x, S_c))] \quad (1)$$

The validation module computes a moving average of this metric over incoming data batches. Concept drift is formally flagged when the adequacy score degrades below a dynamically or statically defined threshold τ (i.e., $A(S_c, D_t) < \tau$). This serves as the trigger mechanism for our adaptation protocol [9].

3. Localized Fine-Tuning Mechanism

Upon detecting that a specific concept representation S_c has decayed, the system initiates a surgical update. Let

$\theta(S_c)$ represent the specific weights and biases directly feeding into and out of the neurons in S_c . Let $\theta(S_c)$ represent all other parameters in the network [9]. The system locks θ_{frozen} (setting their gradient requirements to zero) and solves the following constrained optimization problem using gradient descent:

$$\min_{\theta(S_c)} \mathbb{E}_{(x,y) \in D_t} [L(M(x; \theta_{\text{frozen}}, \theta(S_c)), y)] \quad (2)$$

By restricting the backward pass to only the isolated conceptual sub-network, the algorithm ensures that the model learns the new variations of the concept (e.g., adapting the "cat" concept to include "Sphynx cats") without disrupting the stable feature extractors in the rest of the network [10].

System Architecture and Flowcharts

The proposed framework is designed as a modular pipeline integrated into the standard inference loop of a Deep Neural Network. Unlike static models that operate in a "deploy and forget" manner, our architecture introduces a continuous feedback loop that monitors the

semantic integrity of the model's internal representations [11]. The architecture is divided into two primary functional components: the **Validation Module**, which serves as the diagnostic unit, and the **Targeted Adaptation Module**, which serves as the therapeutic unit.

The operational lifecycle of the system begins with the identification of neural correlates via L_0 -optimization. This process isolates the specific topological regions of the network (neurons and their corresponding weights) that contribute most significantly to a target concept. Once these regions are mapped, the model enters the monitoring phase [12]. In this phase, every batch of incoming data is processed not only for inference but also for adequacy assessment.

The overall logic of the system's decision-making process – from the initial data ingestion to the final weight update – is illustrated in the flowchart below Fig. 1.

As shown in Figure 1, the system relies on the **Concept Adequacy Score** as a gatekeeper. If the score remains




Fig. 1. Flowchart of the dynamic validation and adaptation algorithm




Fig. 2. Structural isolation of parameters during the Targeted Fine-Tuning phase

above the threshold τ , the model continues to operate without intervention, ensuring zero additional training overhead during stable data periods. However, the moment a distribution shift (concept drift) is detected, the workflow branches into the **Targeted Fine-Tuning** partition [13].

The technical novelty of this architecture lies in its ability to perform "surgical" updates. Traditional fine-tuning protocols usually require the adjustment of all weights in the final layers or the entire network. Our approach, however, implements a strict parameter-masking strategy [14]. This structural isolation ensures that the gradients generated during backpropagation are mathematically constrained to the subset of parameters identified during the localization phase.

The physical distribution of these parameters and the state of the network during a localized update are visualized in the structural diagram below [15].

In Figure 2, the distinction between the "frozen" layers and the "active" neurons is clearly defined. By freezing the vast majority of the network's parameters (θ_{frozen}), we effectively lock the high-level feature extractors and historical knowledge in place. Only the specific neural correlates (θ_{active}) associated with the drifting concept are permitted to reach a trainable state.

This architectural design solves the dual challenge of **Catastrophic Forgetting** and **Computational Inefficiency**. Because the update is localized, the "catastrophic" interference with other learned classes is physically impossible, as their corresponding weights remain unchanged [16]. Furthermore, by calculating gradients for only a fraction of the network, the computational demand is reduced to a small percentage of a full training pass, making the system suitable for real-time adaptation in edge computing environments.

DISCUSSION OF RESULTS

The framework was evaluated using a pre-trained VGG16 architecture in a simulated continual learning environment with three distinct stages: Initial Training

(Persian cats), Drift 1 (introduction of new sub-classes like Sphynx), and Drift 2 (introduction of visually similar distinct classes like small dogs) [17].

The experimental results (referencing Table 1) reveal stark contrasts between the compared methodologies. As the data stream evolved, the **Static Model** suffered catastrophic performance drops, degrading from 98.1% to 51.2% by Stage 3, underscoring the severe impact of unmitigated concept drift.

The **Full Retraining** approach maintained the highest accuracy boundaries (96.5% at Stage 3) but incurred a massive computational penalty of approximately 15.5 GFLOPS per update stage. This renders it highly inefficient for continuous deployment [18].

The **Proposed Framework** utilizing localized updates demonstrated exceptional resilience and efficiency. It maintained a highly competitive accuracy profile (96.2% at Stage 2 and 93.1% at Stage 3), sacrificing only ~3.4% accuracy compared to the brute-force full retraining method [19]. However, this marginal trade-off in accuracy yielded profound computational dividends [20]. By executing backpropagation exclusively on the isolated parameter subset $\Theta(S_c)$, the computational cost was reduced to ~0.8 GFLOPS per stage [21]. This represents a staggering **~94.8% reduction in computational overhead** [22].

These results confirm that catastrophic forgetting can be largely bypassed by surgically updating outdated neural correlates, proving that global network updates are not a strict requirement for adapting to localized semantic drift [23].

CONCLUSIONS

This research successfully introduces a paradigm-shifting framework that addresses the vulnerabilities of deep neural networks in non-stationary environments. By synthesizing principles of Explainable AI and Continual Learning, we have transformed static neural correlates into dynamic, maintainable functional units. The primary

contributions of this work are threefold. First, we formalized concept drift at the micro-architectural level by introducing the "Concept Adequacy Score," enabling continuous, automated monitoring of specific knowledge representations. Second, we developed a targeted fine-tuning mechanism that isolates and updates only the necessary mathematical subset of the network. Third, empirical validation proved that this approach drastically reduces the computational burden of model adaptation—slashing GFLOPS requirements by over 90% while maintaining

accuracy within a highly acceptable margin compared to full retraining baselines.

Ultimately, this localized approach provides a highly scalable, computationally lightweight, and interpretable pathway for deploying adaptive Artificial Intelligence systems capable of thriving in real-world, ever-changing conditions. Future research will focus on developing dynamic, self-adjusting thresholds (τ) for the Adequacy Score and expanding the framework to Transformer-based architectures.

REFERENCES

- [1] Liang, Y.-S., & Li, W.-J. (2024). InfLoRA: Interference-Free Low-Rank Adaptation for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1-10.
- [2] Liu, J., Ji, Z., Yu, Y., Cao, J., Pang, Y., Han, J., & Li, X. (2024). Parameter-Efficient Fine-Tuning for Continual Learning: A Neural Tangent Kernel Perspective. *arXiv preprint arXiv:2407.17120*.
- [3] Zhao, L., Zhang, X., Yan, K., Ding, S., & Huang, W. (2024). SAFE: Slow and Fast Parameter-Efficient Tuning for Continual Learning with Pre-Trained Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [4] Zhao, W., Wang, S., Hu, Y., Zhao, Y., Qin, B., Zhang, X., Yang, Q., Xu, D., & Che, W. (2024). SAPT: A Shared Attention Framework for Parameter-Efficient Continual Learning of Large Language Models. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- [5] Yu, J., Zhuge, Y., Zhang, L., Hu, P., Wang, D., Lu, H., & He, Y. (2024). Boosting Continual Learning of Vision-Language Models via Mixture-of-Experts Adapters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [6] D'Alessandro, M., Alonso, A., Calabrés, E., & Galar, M. (2023). Multimodal Parameter-Efficient Few-Shot Class Incremental Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 3385-3395.
- [7] Rymarczyk, D., van de Weijer, J., Zieliński, B., & Twardowski, B. (2023). ICICLE: Interpretable Class Incremental Continual Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [8] Wang, L., Xie, J., Zhang, X., Su, H., & Zhu, J. (2023). Towards a General Framework for Continual Learning with Pre-training. *arXiv preprint arXiv:2310.13888*.
- [9] Yoon, J., Hwang, S. J., & Cao, Y. (2023). Continual Learners are Incremental Model Generalizers. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [10] Tang, C. I., Qendro, L., Spathis, D., Kawsar, F., Mascolo, C., & Mathur, A. (2024). Kaizen: Practical Self-supervised Continual Learning with Continual Fine-tuning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2841-2850.
- [11] Wang, Y., et al. (2024). InsCL: A Data-efficient Continual Learning Paradigm for Fine-tuning Large Language Models with Instructions. *arXiv preprint arXiv:2403.11435*.
- [12] Tang, C. I., Qendro, L., Spathis, D., Kawsar, F., Mathur, A., & Mascolo, C. (2024). Balancing Continual Learning and Fine-tuning for Human Activity Recognition. In *AAAI 2024 HCRL Workshop*.
- [13] Wistuba, M., et al. (2023). Continual Learning with Low Rank Adaptation. In *NeurIPS 2023 Workshop on Distribution Shifts*.
- [14] Song, C., et al. (2023). ConPET: Continual Parameter-Efficient Tuning for Large Language Models. *arXiv preprint arXiv:2309.14763*.
- [15] Chen, Y., et al. (2023). Parameterizing Context: Unleashing the Power of Parameter-Efficient Fine-Tuning and In-Context Tuning for Continual Table Semantic Parsing. *arXiv preprint arXiv:2310.04801*.
- [16] Weyssow, M., et al. (2023). On the Usage of Continual Learning for Out-of-Distribution Generalization in Pre-trained Language Models of Code. *arXiv preprint arXiv:2305.04106*.
- [17] Jiao, J., et al. (2025). DriftMoE: A Mixture of Experts Approach to Handle Concept Drifts. *arXiv preprint arXiv:2507.18464*.
- [18] Yu, S., et al. (2025). Holistic Continual Learning under Concept Drift with Adaptive Memory Realignment. *arXiv preprint arXiv:2507.02310*.
- [19] Lu, J., et al. (2025). Evolving Machine Learning in Non-Stationary Environments: A Unified Survey of Drift, Forgetting, and Adaptation. *arXiv preprint arXiv:2505.17902*.
- [20] Liu, Z., et al. (2025). Parameter-Efficient Continual Fine-Tuning: A Survey. *arXiv preprint arXiv:2504.13822*.
- [21] Xu, M., et al. (2025). Continual Learning in Vision-Language Models via Aligned Model Merging. *arXiv preprint arXiv:2506.03189*.
- [22] Yang, C., Lu, H., & Yu, P. (2025). Drift-aware Collaborative Assistance Mixture of Experts for Heterogeneous Multistream Learning. *arXiv preprint arXiv:2508.01598*.
- [23] Wan, L., et al. (2024). Proactive Model Adaptation Against Concept Drift for Online Time Series Forecasting. *arXiv preprint arXiv:2412.08435*.

СПИСОК ВИКОРИСТАНОЇ ЛІТЕРАТУРИ

- [1] Liang, Y.-S., & Li, W.-J. (2024). InfLoRA: Interference-Free Low-Rank Adaptation for Continual Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1-10.
- [2] Liu, J., Ji, Z., Yu, Y., Cao, J., Pang, Y., Han, J., & Li, X. (2024). Parameter-Efficient Fine-Tuning for Continual Learning: A Neural Tangent Kernel Perspective. *arXiv preprint arXiv:2407.17120*.
- [3] Zhao, L., Zhang, X., Yan, K., Ding, S., & Huang, W. (2024). SAFE: Slow and Fast Parameter-Efficient Tuning for Continual Learning with Pre-Trained Models. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [4] Zhao, W., Wang, S., Hu, Y., Zhao, Y., Qin, B., Zhang, X., Yang, Q., Xu, D., & Che, W. (2024). SAPT: A Shared Attention Framework for Parameter-Efficient Continual Learning of Large Language Models. In *Proceedings of the Association for Computational Linguistics (ACL)*.
- [5] Yu, J., Zhuge, Y., Zhang, L., Hu, P., Wang, D., Lu, H., & He, Y. (2024). Boosting Continual Learning of Vision-Language Models via Mixture-of-Experts Adapters. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [6] D'Alessandro, M., Alonso, A., Calabrés, E., & Galar, M. (2023). Multimodal Parameter-Efficient Few-Shot Class Incremental Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 3385-3395.
- [7] Rymarczyk, D., van de Weijer, J., Zieliński, B., & Twardowski, B. (2023). ICICLE: Interpretable Class Incremental Continual Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- [8] Wang, L., Xie, J., Zhang, X., Su, H., & Zhu, J. (2023). Towards a General Framework for Continual Learning with Pre-training. *arXiv preprint arXiv:2310.13888*.
- [9] Yoon, J., Hwang, S. J., & Cao, Y. (2023). Continual Learners are Incremental Model Generalizers. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [10] Tang, C. I., Qendro, L., Spathis, D., Kawsar, F., Mascolo, C., & Mathur, A. (2024). Kaizen: Practical Self-supervised Continual Learning with Continual Fine-tuning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2841-2850.
- [11] Wang, Y., et al. (2024). InsCL: A Data-efficient Continual Learning Paradigm for Fine-tuning Large Language Models with Instructions. *arXiv preprint arXiv:2403.11435*.
- [12] Tang, C. I., Qendro, L., Spathis, D., Kawsar, F., Mathur, A., & Mascolo, C. (2024). Balancing Continual Learning and Fine-tuning for Human Activity Recognition. In *AAAI 2024 HCRL Workshop*.
- [13] Wistuba, M., et al. (2023). Continual Learning with Low Rank Adaptation. In *NeurIPS 2023 Workshop on Distribution Shifts*.
- [14] Song, C., et al. (2023). ConPET: Continual Parameter-Efficient Tuning for Large Language Models. *arXiv preprint arXiv:2309.14763*.
- [15] Chen, Y., et al. (2023). Parameterizing Context: Unleashing the Power of Parameter-Efficient Fine-Tuning and In-Context Tuning for Continual Table Semantic Parsing. *arXiv preprint arXiv:2310.04801*.
- [16] Weyssow, M., et al. (2023). On the Usage of Continual Learning for Out-of-Distribution Generalization in Pre-trained Language Models of Code. *arXiv preprint arXiv:2305.04106*.
- [17] Jiao, J., et al. (2025). DriftMoE: A Mixture of Experts Approach to Handle Concept Drifts. *arXiv preprint arXiv:2507.18464*.
- [18] Yu, S., et al. (2025). Holistic Continual Learning under Concept Drift with Adaptive Memory Realignment. *arXiv preprint arXiv:2507.02310*.
- [19] Lu, J., et al. (2025). Evolving Machine Learning in Non-Stationary Environments: A Unified Survey of Drift, Forgetting, and Adaptation. *arXiv preprint arXiv:2505.17902*.
- [20] Liu, Z., et al. (2025). Parameter-Efficient Continual Fine-Tuning: A Survey. *arXiv preprint arXiv:2504.13822*.
- [21] Xu, M., et al. (2025). Continual Learning in Vision-Language Models via Aligned Model Merging. *arXiv preprint arXiv:2506.03189*.
- [22] Yang, C., Lu, H., & Yu, P. (2025). Drift-aware Collaborative Assistance Mixture of Experts for Heterogeneous Multistream Learning. *arXiv preprint arXiv:2508.01598*.
- [23] Wan, L., et al. (2024). Proactive Model Adaptation Against Concept Drift for Online Time Series Forecasting. *arXiv preprint arXiv:2412.08435*.

© Вербицький О. С.

Дата першого надходження статті до видання: 23.02.2026

Дата прийняття статті до друку після рецензування: 19.03.2026

Дата публікації (оприлюднення) статті: 22.05.2026



Стаття поширюється на умовах ліцензії відкритого доступу (CC BY 4.0)