Preprints.org

Review

Continual Learning in Artificial Intelligence: A Review of Techniques, Metrics, and Real-World Applications

Nereo Taygete , Stephanie Waldemar , Husniya Salwa [*]

Review

Continual Learning in Artificial Intelligence: A Review of Techniques, Metrics, and Real-World Applications

Nereo Taygete [1], Stephanie Waldemar [2] and Husniya Salwa [1,2,*]

[1] King Abdullah University of Science and Technology, Saudi Arabia
[2] King Abdulaziz University, Saudi Arabia
[*] Correspondence: HusniyaSalwa@kaust.edu.sa

Abstract: Continual learning (CL) is a critical paradigm in artificial intelligence that enables models to learn sequentially from a stream of tasks while retaining previously acquired knowledge. Unlike traditional machine learning approaches that assume static datasets, CL aims to address real-world scenarios where data distributions evolve over time. However, CL models face significant challenges, including catastrophic forgetting, scalability, task interference, and the trade-off between stability and plasticity. This survey provides a comprehensive review of continual learning, covering key learning paradigms such as regularization-based methods, memory replay techniques, and dynamic architectural approaches. We discuss widely used evaluation metrics, benchmark datasets, and experimental protocols that facilitate fair comparisons among CL methods. Additionally, we explore real-world applications of CL in domains such as robotics, healthcare, natural language processing, cybersecurity, and recommender systems. Despite recent advances, several open challenges remain, including efficient memory management, task-free learning, privacy concerns, and improving forward and backward transfer. We highlight emerging research directions, including neuroscience-inspired learning mechanisms, self-supervised continual learning, meta-learning, and multi-modal CL. Finally, we discuss the integration of CL into large-scale foundation models and human-AI collaborative systems. By presenting an in-depth analysis of continual learning methodologies, challenges, and future prospects, this survey aims to provide researchers and practitioners with a structured understanding of the field and inspire further innovations in building adaptive, lifelong learning AI systems.

Keywords: continual learning; catastrophic forgetting; lifelong learning; memory replay; task-incremental learning; evaluation metrics; applications; scalability; neural networks

1. Introduction

Machine learning has made significant strides in recent years, achieving state-of-the-art performance in various domains, including computer vision, natural language processing, robotics, and healthcare [1]. These successes, however, have largely been achieved under the assumption that models are trained on static datasets and evaluated on similar data distributions. In real-world scenarios, data distributions often evolve, new tasks emerge over time, and models must continuously integrate new knowledge while retaining previously learned information. Traditional machine learning models, when trained sequentially on new data, suffer from *catastrophic forgetting*, where the performance on previously learned tasks deteriorates as new information overwrites prior knowledge. Continual learning (CL), also known as lifelong learning or incremental learning, aims to address these challenges by enabling models to learn in a sequential manner while maintaining stability and adaptability [2]. Unlike conventional machine learning approaches that require retraining from scratch when new data becomes available, continual learning seeks to develop models that can learn continuously from new experiences while preserving past knowledge. Achieving this goal requires overcoming key challenges such as mitigating catastrophic forgetting, efficiently leveraging computational and memory resources, and generalizing across diverse and evolving tasks. Several strategies have been proposed in the

literature to tackle the problem of continual learning. These approaches can be broadly categorized into three major paradigms:

- **Replay-based methods:** These techniques mitigate forgetting by storing past data explicitly in memory buffers or generating synthetic samples to rehearse previous knowledge.
- **Regularization-based methods:** Regularization techniques impose constraints on model updates to retain previously learned knowledge, often by penalizing drastic changes in important parameters [3].
- **Dynamic architecture-based methods:** These approaches dynamically expand or modify the model architecture to accommodate new tasks while preserving existing knowledge.

The evaluation of continual learning models remains a crucial research challenge, as standard machine learning benchmarks do not adequately capture the complexities of sequential learning [4]. Metrics such as *accuracy drop*, *forward transfer*, and *backward transfer* have been proposed to assess the effectiveness of CL methods. Additionally, benchmark datasets such as *Permuted MNIST*, *Split CIFAR*, and *CORe50* have been introduced to facilitate comparative analysis of different CL approaches [5]. Continual learning has broad implications for real-world applications, including autonomous systems, personalized recommendation engines, and adaptive healthcare models [6]. In robotics, for instance, continual learning enables intelligent agents to incrementally acquire new skills without forgetting previously learned behaviors. In healthcare, adaptive diagnostic models can improve their accuracy over time as new patient data becomes available, reducing the need for costly retraining [7]. This survey provides a comprehensive overview of continual learning by reviewing key methodologies, challenges, and state-of-the-art advancements in the field. We categorize and analyze existing approaches, discuss evaluation metrics and benchmark datasets, and highlight real-world applications [8]. Furthermore, we explore open research challenges and future directions, aiming to provide valuable insights for researchers and practitioners working on continual learning. Through this survey, we seek to foster further research and innovation

2. Background and Problem Definition

Continual learning (CL) represents a fundamental shift from traditional machine learning paradigms, which assume that models are trained on static datasets and evaluated in similar environments [9]. Instead, CL models are designed to learn sequentially from a continuous stream of data, adapting to new tasks while retaining previously acquired knowledge [10]. This section provides a formal definition of the continual learning problem and discusses the key challenges associated with learning in a dynamic setting [11].

2.1. Formal Definition of Continual Learning

In a continual learning setup, a model is trained on a sequence of tasks $T_1, T_2, \ldots, T_N$, where each task T_i is associated with a dataset $D_i = \{(x_j, y_j)\}_{j=1}^{N_i}$ drawn from a data distribution $P_i(X, Y)$ [12]. The goal of continual learning is to train a model f_θ with parameters θ that can learn from each new task T_i while maintaining high performance on previously learned tasks $T_1, T_2, \ldots, T_{i-1}$ [13]. Mathematically, a continual learning model should aim to minimize the overall loss:

$$\mathcal{L}_{\text{CL}} = \sum_{i=1}^{N} \mathbb{E}_{(x,y) \sim P_i}[\mathcal{L}(f_\theta(x), y)] \tag{1}$$

where $\mathcal{L}$ is the task-specific loss function (e.g., cross-entropy loss for classification tasks) [14]. The primary challenge in continual learning is to update the model on new tasks without suffering from catastrophic forgetting, where performance on earlier tasks deteriorates due to the sequential learning process [15].

2.2. Challenges in Continual Learning

Several challenges arise in designing effective continual learning algorithms:

2.2.1. Catastrophic Forgetting

A major issue in CL is *catastrophic forgetting*, where a model, upon learning new tasks, experiences a severe degradation in performance on previously learned tasks. This occurs because standard gradient-based optimization updates model parameters based on new task data, potentially overwriting previously learned information [16].

2.2.2. Knowledge Transfer and Interference

Continual learning models must balance knowledge transfer and interference [17]. Ideally, previously learned knowledge should be leveraged to improve learning on new tasks (*positive transfer*), while avoiding interference where learning new tasks negatively impacts performance on old ones [18].

2.2.3. Resource Constraints

Real-world continual learning settings often impose constraints on memory, computation, and data storage. Unlike traditional machine learning, where a full dataset is available for training, CL systems may have limited or no access to prior task data, making efficient learning strategies essential.

2.2.4. Task Boundaries and Learning Paradigms

Continual learning can be categorized into different learning paradigms:

- **Task-incremental learning:** Task identities are known, and separate task-specific classifiers may be used [19].
- **Domain-incremental learning:** The same task is learned under shifting data distributions [20].
- **Class-incremental learning:** New classes are introduced over time, and the model must integrate them into a unified classifier.

Each paradigm presents unique challenges in terms of model design, adaptation, and evaluation.

2.3. Comparison with Related Learning Paradigms

Continual learning is closely related to other machine learning paradigms, but with key differences:

- **Online Learning:** Online learning processes data in a sequential manner but does not necessarily retain knowledge from past data distributions, whereas continual learning aims to accumulate knowledge over time.
- **Meta-Learning:** Meta-learning focuses on learning how to learn across multiple tasks, whereas continual learning focuses on long-term retention and adaptation [21].
- **Multi-Task Learning:** Multi-task learning trains a model on multiple tasks simultaneously, whereas continual learning handles tasks sequentially [22].

2.4. Importance of Continual Learning

Continual learning is crucial for developing adaptive, intelligent systems capable of operating in dynamic environments. Its applications span various domains, including:

- **Autonomous Systems:** Self-driving cars and robotic agents must continuously learn from new interactions and environments.
- **Healthcare:** Diagnostic models must adapt to evolving medical data without retraining from scratch [23].
- **Personalized AI:** User-adaptive AI systems, such as recommendation engines and virtual assistants, require continual adaptation based on user preferences.

As continual learning continues to evolve, addressing these challenges will be crucial to advancing AI systems that can learn in a lifelong manner [24]. In the following sections, we provide a detailed survey of the various approaches developed to tackle these challenges.

3. Taxonomy of Continual Learning Approaches

Continual learning (CL) has been tackled using various approaches, each aiming to address the challenges of catastrophic forgetting, efficient resource utilization, and effective knowledge transfer. Broadly, these approaches can be categorized into three main families: (1) replay-based methods, (2) regularization-based methods, and (3) dynamic architecture-based methods. In this section, we provide an in-depth discussion of each category, highlighting key techniques and representative works [25].

3.1. Replay-Based Methods

Replay-based methods mitigate forgetting by storing past experiences or generating synthetic data to reinforce previously learned knowledge [26]. These methods leverage explicit memory mechanisms to maintain performance across sequential tasks.

3.1.1. Experience Replay

Experience replay involves storing a subset of past samples in a memory buffer and periodically replaying them alongside new data during training. This approach simulates the presence of past tasks, reducing the risk of catastrophic forgetting.

- **iCaRL** (Incremental Classifier and Representation Learning) maintains a memory buffer of exemplars and uses nearest-neighbor classification to adapt to new tasks [27].
- **AGEM** (Average Gradient Episodic Memory) constrains gradient updates using stored samples to prevent drastic parameter changes [28].

3.1.2. Generative Replay

Instead of storing raw data, generative replay techniques use generative models such as Variational Autoencoders (VAEs) or Generative Adversarial Networks (GANs) to synthesize past data.

- **DGR** (Deep Generative Replay) trains a generative model to generate previous task data, which is then replayed alongside new task data.
- **Brain-Inspired Replay** mimics biological memory consolidation by incorporating generative models for long-term knowledge retention.

3.2. Regularization-Based Methods

Regularization-based methods aim to prevent catastrophic forgetting by introducing additional loss terms that constrain weight updates, ensuring that previously learned knowledge is preserved.

3.2.1. Penalty-Based Regularization

These methods add penalty terms to the loss function to prevent drastic changes in important model parameters.

- **Elastic Weight Consolidation (EWC)** estimates the importance of each parameter using Fisher information and penalizes changes to important weights [29].
- **Synaptic Intelligence (SI)** tracks the influence of parameters during training and restricts significant modifications [30].

3.2.2. Knowledge Distillation

Knowledge distillation techniques encourage consistency between the outputs of a current model and a previously trained model to retain past knowledge [31].

- **Learning without Forgetting (LwF)** uses distillation loss to ensure that the model's predictions on old tasks remain stable while learning new tasks.
- **Variational Continual Learning (VCL)** integrates Bayesian learning with knowledge distillation to improve stability.

3.3. Dynamic Architecture-Based Methods

Dynamic architecture-based methods address continual learning by expanding, modifying, or selectively freezing parts of the network to accommodate new tasks while preserving prior knowledge.

3.3.1. Network Expansion

These methods allocate additional network capacity when learning new tasks to minimize interference with previous knowledge [32].

- **Progressive Neural Networks (PNN)** create a new sub-network for each task while maintaining lateral connections to prior networks.
- **Dynamically Expandable Networks (DEN)** selectively grow network components while reusing prior representations [33].

3.3.2. Parameter Isolation

Parameter isolation techniques assign specific subnetworks to different tasks, reducing interference between them [34].

- **PathNet** enables task-specific routing by selecting a subset of network paths for each task.
- **Supermask Superposition** learns task-specific masks over a fixed network to enable multi-task learning without forgetting [35].

3.4. Hybrid Approaches

Several methods integrate multiple CL strategies to achieve better performance. For example:

- **MER (Meta-Experience Replay)** combines experience replay with meta-learning to enhance sample efficiency [36].
- **ER-RingBuffer** uses a memory buffer with a regularization mechanism to balance stability and adaptability.

3.5. Comparison of Approaches

Each category of continual learning methods has its strengths and limitations:

- Replay-based methods provide strong performance but require additional memory storage [37].
- Regularization-based methods are memory-efficient but may struggle with long-term retention [38].
- Dynamic architecture methods offer flexibility but can lead to scalability challenges.

Understanding these trade-offs is crucial for selecting the right approach based on application requirements and constraints.

4. Evaluation Metrics and Benchmarks

Evaluating continual learning (CL) models presents unique challenges compared to traditional machine learning due to the sequential nature of learning and the risk of catastrophic forgetting [39]. To ensure fair comparisons, researchers rely on specific evaluation metrics and benchmark datasets [40]. This section details the most commonly used evaluation metrics and benchmark datasets in continual learning [41].

4.1. Evaluation Metrics

A well-designed continual learning model should effectively retain past knowledge, transfer knowledge across tasks, and efficiently adapt to new data. The following metrics are commonly used to assess these capabilities:

4.1.1. Average Accuracy (ACC)

The average accuracy across all learned tasks provides a measure of overall model performance [42]. Given N tasks, the accuracy on task T_i at the end of training is denoted as $A_{i,N}$. The average accuracy is computed as:

$$ACC = \frac{1}{N} \sum_{i=1}^{N} A_{i,N} \tag{2}$$

Higher ACC values indicate better knowledge retention and generalization across tasks.

4.1.2. Forgetting Measure (FM)

The forgetting measure quantifies how much performance on previous tasks degrades as new tasks are learned. It is defined as:

$$FM = \frac{1}{N-1} \sum_{i=1}^{N-1} \max_{j \in \{1,\ldots,N-1\}} (A_{i,j} - A_{i,N}) \tag{3}$$

where $A_{i,j}$ represents the accuracy of task i after learning task j. A lower FM indicates less forgetting [43].

4.1.3. Forward Transfer (FWT)

Forward transfer measures the improvement in learning new tasks due to knowledge gained from previous tasks [44]. It is defined as:

$$FWT = \frac{1}{N-1} \sum_{i=2}^{N} (A_{i,i} - A_{i,0}) \tag{4}$$

where $A_{i,0}$ is the accuracy of task i if trained from scratch. Higher FWT values indicate better transfer of knowledge to new tasks [45].

4.1.4. Backward Transfer (BWT)

Backward transfer measures the effect of learning new tasks on previously learned tasks:

$$BWT = \frac{1}{N-1} \sum_{i=1}^{N-1} (A_{i,N} - A_{i,i}) \tag{5}$$

A positive BWT suggests that learning new tasks improves performance on previous tasks, while a negative BWT indicates forgetting [46].

4.1.5. Memory Overhead

Since continual learning methods often rely on additional memory buffers, it is important to assess the memory footprint. Methods such as experience replay require storing past data, whereas dynamic architectures may require additional parameters [47]. Memory efficiency is evaluated based on the percentage of additional storage required compared to a standard model [48–50].

4.2. Benchmark Datasets

To fairly compare CL algorithms, standard benchmark datasets with well-defined task sequences are commonly used. Below are some widely adopted datasets:

4.2.1. Permuted MNIST

Permuted MNIST is a synthetic benchmark where each task consists of a permuted version of the MNIST dataset. It evaluates a model's ability to generalize across distorted input distributions [51].

4.2.2. Split MNIST

Split MNIST divides the ten digits of MNIST into multiple binary classification tasks (e.g., distinguishing 0s from 1s, 2s from 3s) [52]. This setup tests a model's ability to incrementally learn new classes.

4.2.3. Split CIFAR-10 and Split CIFAR-100

Split CIFAR-10 and Split CIFAR-100 divide image classes into multiple subsets, forming a sequence of tasks. These benchmarks are used to assess class-incremental learning performance.

4.2.4. CORe50

CORe50 is a continual object recognition dataset consisting of 50 objects captured under different conditions [53]. It is commonly used for evaluating CL in real-world scenarios.

4.2.5. TinyImageNet

TinyImageNet is a smaller version of ImageNet, often used in continual learning settings where models incrementally learn new categories.

4.2.6. Omniglot

Omniglot consists of handwritten characters from multiple alphabets and is used to study continual learning in few-shot settings.

4.3. Experimental Protocols

Different experimental setups are used to assess continual learning performance:

- **Task-Incremental Learning (Task-IL):** The model is provided with task identifiers and learns separate classifiers for each task.
- **Domain-Incremental Learning (Domain-IL):** The model encounters data from the same set of classes but under different distributions (e.g., lighting changes in images).
- **Class-Incremental Learning (Class-IL):** New classes are introduced over time, and the model must integrate them into a single classifier [54].

By evaluating continual learning models across these metrics, datasets, and experimental protocols, researchers can gain insights into their effectiveness and limitations [55]. The next section explores real-world applications of continual learning in various domains [56].

5. Applications of Continual Learning

Continual learning (CL) has emerged as a crucial capability for intelligent systems operating in dynamic environments [57]. The ability to learn incrementally without forgetting previous knowledge is essential in numerous real-world applications [58]. In this section, we explore key domains where continual learning has demonstrated significant impact [59].

5.1. Autonomous Systems and Robotics

Autonomous agents, such as self-driving cars and robots, must continuously adapt to changing environments, new tasks, and unexpected conditions [60]. Continual learning enables these systems to improve their perception, decision-making, and control over time [61].

5.1.1. Self-Driving Cars

Self-driving vehicles operate in highly dynamic environments where new road conditions, traffic patterns, and obstacles emerge regularly [62]. Continual learning allows models to:

- Adapt to new weather conditions and road structures without retraining from scratch.
- Improve object detection and classification as new obstacles or traffic signs appear.
- Learn from real-time driving experiences to enhance trajectory planning and collision avoidance.

5.1.2. Robotics and Human-Robot Interaction

Robots deployed in industrial settings, healthcare, and homes must continuously acquire new skills and adapt to different users or environments. Continual learning helps in:

- Learning new manipulation tasks without forgetting previously learned ones [63].
- Enhancing speech and gesture recognition for improved human-robot communication [64].
- Adapting to novel objects and environments in real-time [65].

5.2. Healthcare and Medical Diagnosis

The medical field is an ever-evolving domain where continual learning can support adaptive diagnostics, personalized treatments, and efficient medical imaging analysis [66].

5.2.1. Medical Imaging and Diagnosis

Medical imaging models must adapt to new diseases, imaging techniques, and patient demographics [67]. Continual learning enables:

- Incremental learning of new disease patterns (e.g., emerging virus strains).
- Adaptation to different medical imaging devices and data distributions [68].
- Reducing the need for large-scale retraining, which is costly and time-consuming.

5.2.2. Personalized Medicine

Healthcare treatments are becoming increasingly personalized based on patient-specific data [69]. Continual learning helps in:

- Adapting treatment recommendations based on patient history [70].
- Learning from patient responses to therapies to refine predictive models [71].
- Handling evolving medical knowledge by integrating new clinical research findings.

5.3. Natural Language Processing (NLP)

Modern NLP models require adaptation to changing language patterns, user preferences, and new vocabulary. Continual learning plays a critical role in:

- Enhancing virtual assistants and chatbots by learning from ongoing interactions.
- Improving language translation systems by integrating new phrases and regional dialects [72].
- Reducing catastrophic forgetting in multi-lingual and domain-adaptive NLP models [73].

5.4. Finance and Fraud Detection

Financial institutions must continuously adapt to evolving fraud patterns, market trends, and regulatory changes [74]. Continual learning aids in:

- Detecting emerging fraudulent activities by updating fraud detection models in real-time [75].
- Adapting risk assessment models to new financial behaviors and global events [76].
- Enhancing algorithmic trading by learning from changing market conditions [77].

5.5. Cybersecurity and Threat Detection

Cybersecurity threats evolve rapidly, requiring adaptive defense mechanisms. Continual learning enhances:

- Intrusion detection systems that recognize new attack patterns while retaining knowledge of previous threats [78].
- Malware classification models that adapt to evolving cyber threats [79].
- Network security systems that update policies dynamically based on emerging vulnerabilities.

5.6. Recommender Systems and Personalized AI

Recommendation engines for e-commerce, streaming services, and social media must adapt to user preferences that change over time. Continual learning enables:

- Personalized content recommendations that evolve with user behavior [80].
- Dynamic adaptation to new product trends and emerging market preferences.
- Real-time updates to user preference models for better engagement.

5.7. Scientific Discovery and Research

In fields like physics, biology, and climate science, continual learning can accelerate scientific progress by:

- Enabling adaptive models that incorporate new experimental data [81].
- Improving climate models by continuously integrating new observations [82].
- Assisting in drug discovery by learning from evolving biochemical interactions [83].

5.8. Edge AI and IoT Devices

Edge AI and Internet of Things (IoT) devices require lightweight continual learning approaches to operate efficiently in resource-constrained environments. Applications include:

- Smart home devices that adapt to user preferences over time [84].
- Industrial IoT systems that optimize performance based on sensor data.
- Wearable health monitoring devices that learn from user activity and vitals.

5.9. Summary

Continual learning has far-reaching applications across multiple domains, from autonomous systems to healthcare and cybersecurity [85]. Its ability to facilitate lifelong learning, adaptation, and knowledge retention makes it a vital component of future AI-driven systems. The next section discusses open challenges and future directions in continual learning research [86].

6. Challenges and Future Directions

Despite significant advancements in continual learning (CL), numerous challenges remain. Addressing these challenges is crucial for building robust and efficient CL systems that can adapt seamlessly to new tasks while preserving past knowledge. In this section, we discuss key challenges and outline promising future research directions.

6.1. Challenges in Continual Learning

6.1.1. Catastrophic Forgetting

One of the most persistent challenges in CL is catastrophic forgetting, where learning new tasks leads to the degradation of previously acquired knowledge [87]. While approaches such as replay, regularization, and dynamic architectures help mitigate this issue, achieving truly stable learning without interference remains an open problem [88].

6.1.2. Scalability and Computational Efficiency

Most CL methods require additional storage (e.g., memory buffers, expanded architectures) or computational overhead (e.g., regularization constraints, generative models)[89]. Ensuring scalability in large-scale applications with high-dimensional data and long task sequences is a critical challenge [90,91].

6.1.3. Task-Free and Online Learning

Many existing CL approaches assume task boundaries are predefined and that training is conducted in separate task phases. However, real-world scenarios often involve continuous data streams

with no explicit task demarcation. Developing task-free and online CL methods that dynamically learn without explicit task labels is an important research direction.

6.1.4. Transfer Learning vs. Interference

Effective continual learning requires balancing knowledge transfer and minimizing negative interference. While forward and backward transfer can improve learning efficiency, poorly managed transfer can degrade performance [92]. Understanding how to selectively transfer knowledge remains a fundamental challenge.

6.1.5. Evaluation Protocols and Standardization

The lack of standardized evaluation protocols makes it difficult to compare different CL methods fairly [93]. Existing benchmarks vary in complexity, data availability, and assumptions about task structure [94]. Establishing universally accepted CL evaluation frameworks is essential for reproducible research [95].

6.1.6. Memory and Privacy Constraints

Many CL methods rely on storing past data (e.g., experience replay), which may not be feasible in privacy-sensitive applications such as healthcare or personalized AI. Developing memory-efficient and privacy-preserving CL methods remains an open challenge.

6.1.7. The Stability-Plasticity Dilemma

Continual learning models must balance stability (retaining old knowledge) and plasticity (adapting to new knowledge). Overemphasizing stability can hinder learning new tasks, while excessive plasticity increases the risk of forgetting [96]. Finding optimal trade-offs remains a key research problem [97].

6.2. Future Directions in Continual Learning

6.2.1. Neuroscience-Inspired Learning Mechanisms

Biological brains exhibit remarkable continual learning capabilities. Insights from neuroscience, such as synaptic consolidation, memory replay, and hierarchical memory structures, could inspire more robust CL algorithms [98].

6.2.2. Self-Supervised and Unsupervised Continual Learning

Most CL research focuses on supervised learning, but real-world scenarios often involve unlabeled or sparsely labeled data [99]. Exploring self-supervised and unsupervised CL techniques could reduce dependence on labeled data and improve generalization.

6.2.3. Meta-Learning for Continual Learning

Meta-learning (learning to learn) can improve CL by enabling models to quickly adapt to new tasks while mitigating forgetting. Incorporating meta-learning strategies can enhance sample efficiency and transfer learning in CL settings.

6.2.4. Lifelong Multi-Modal Learning

Future AI systems will need to integrate information from multiple modalities (e.g., vision, language, audio) [100]. Developing CL models capable of learning from diverse data sources while maintaining coherence across modalities is an exciting challenge.

6.2.5. Continual Learning for Foundation Models

Large-scale foundation models (e.g., GPT, BERT, CLIP) are typically trained in a static manner. Integrating CL capabilities into such models could enable them to adapt to new knowledge over time without requiring expensive retraining [101].

6.2.6. Human-AI Collaboration in Continual Learning

Interactive learning, where humans provide feedback to guide model adaptation, is a promising area for CL research. Developing human-in-the-loop CL systems can improve transparency, interpretability, and adaptability [102].

6.2.7. Applications in Real-World Dynamic Environments

CL has great potential in applications such as robotics, autonomous systems, and adaptive user interfaces [103]. Future research should focus on deploying CL models in real-world scenarios where environmental changes are continuous and unpredictable.

6.3. Summary

Continual learning remains a rapidly evolving field with many open challenges and exciting research opportunities. Addressing catastrophic forgetting, improving scalability, and developing task-free learning paradigms are critical for advancing the field [104]. By drawing inspiration from neuroscience, exploring self-supervised approaches, and integrating CL into foundation models, future research can enable truly intelligent and adaptive AI systems.

7. Conclusion

Continual learning (CL) is a fundamental challenge in artificial intelligence, aiming to develop models that can learn incrementally while retaining previously acquired knowledge. This survey has provided a comprehensive overview of CL, including its foundational concepts, learning paradigms, evaluation metrics, benchmarks, real-world applications, and open challenges.

We first explored the key principles of continual learning, highlighting different approaches such as regularization-based methods, memory replay techniques, and dynamic architectural adaptations. We then examined various evaluation metrics used to assess CL performance, including average accuracy, forgetting measures, and transfer learning effectiveness. Benchmark datasets, such as Permuted MNIST, Split CIFAR, and CORe50, were discussed as standardized testing grounds for CL models.

Continual learning has demonstrated significant potential in various real-world applications, including autonomous systems, healthcare, natural language processing, cybersecurity, and recommender systems. However, despite its progress, several challenges remain, particularly in mitigating catastrophic forgetting, ensuring scalability, and enabling task-free online learning. Additionally, privacy constraints and computational efficiency pose further obstacles in deploying CL in practical settings.

Looking ahead, future research should focus on integrating neuroscience-inspired mechanisms, advancing self-supervised and meta-learning approaches, and developing multi-modal CL models capable of learning from diverse data streams. Furthermore, improving CL techniques for foundation models and exploring human-AI collaborative learning will be crucial in creating more adaptive and intelligent AI systems.

In conclusion, continual learning remains an active and rapidly evolving research field with far-reaching implications for the future of artificial intelligence. By addressing the existing challenges and embracing new research directions, CL can pave the way for truly lifelong learning systems capable of seamlessly adapting to dynamic environments.

References

1. Yoon, J.; Jeong, W.; Lee, G.; Yang, E.; Hwang, S.J. Federated continual learning with weighted inter-client transfer. In Proceedings of the International Conference on Machine Learning. PMLR, 2021, pp. 12073–12086.
2. Caccia, M.; Rodriguez, P.; Ostapenko, O.; Normandin, F.; Lin, M.; Page-Caccia, L.; Laradji, I.H.; Rish, I.; Lacoste, A.; Vázquez, D.; et al. Online fast adaptation and knowledge accumulation (osaka): a new approach to continual learning. *Advances in Neural Information Processing Systems* **2020**, *33*, 16532–16545.

3. Wang, Z.; Mehta, S.V.; Póczos, B.; Carbonell, J.G. Efficient Meta Lifelong-Learning with Limited Memory. In Proceedings of the Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 535–548.

4. Mendez, J.A.; Eaton, E. A General Framework for Continual Learning of Compositional Structures.

5. Hurtado, J.; Raymond, A.; Soto, A. Optimizing reusable knowledge for continual learning via metalearning. *Advances in Neural Information Processing Systems* **2021**, *34*, 14150–14162.

6. Li, Y.; Zhao, L.; Church, K.; Elhoseiny, M. Compositional Language Continual Learning. In Proceedings of the International Conference on Learning Representations, 2019.

7. Golatkar, A.; Achille, A.; Ravichandran, A.; Polito, M.; Soatto, S. Mixed-privacy forgetting in deep networks. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 792–801.

8. Ehret, B.; Henning, C.; Cervera, M.; Meulemans, A.; Von Oswald, J.; Grewe, B.F. Continual learning in recurrent neural networks. In Proceedings of the International Conference on Learning Representations, 2020.

9. Cermelli, F.; Fontanel, D.; Tavera, A.; Ciccone, M.; Caputo, B. Incremental learning in semantic segmentation from image labels. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 4371–4381.

10. Zhu, F.; Zhang, X.Y.; Wang, C.; Yin, F.; Liu, C.L. Prototype augmentation and self-supervision for incremental learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 5871–5880.

11. Wei, N.; Zhou, T.; Zhang, Z.; Zhuo, Y.; Chen, L. Visual working memory representation as a topological defined perceptual object. *Journal of Vision* **2019**, *19*, 12–12.

12. Douillard, A.; Chen, Y.; Dapogny, A.; Cord, M. Plop: Learning without forgetting for continual semantic segmentation. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 4040–4050.

13. Ke, Z.; Liu, B.; Ma, N.; Xu, H.; Shu, L. Achieving Forgetting Prevention and Knowledge Transfer in Continual Learning. *Advances in Neural Information Processing Systems* **2021**, *34*, 22443–22456.

14. Arevian, A.C.; Kapoor, V.; Urban, N.N. Activity-dependent gating of lateral inhibition in the mouse olfactory bulb. *Nature Neuroscience* **2008**, *11*, 80–87.

15. Singh, P.; Verma, V.K.; Mazumder, P.; Carin, L.; Rai, P. Calibrating cnns for lifelong learning. *Advances in Neural Information Processing Systems* **2020**, *33*, 15579–15590.

16. Aso, Y.; Sitaraman, D.; Ichinose, T.; Kaun, K.R.; Vogt, K.; Belliart-Guérin, G.; Plaçais, P.Y.; Robie, A.A.; Yamagata, N.; Schnaitmann, C.; et al. Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila. *Elife* **2014**, *3*, e04580.

17. Michieli, U.; Zanuttigh, P. Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 1114–1124.

18. Li, J.; Ji, Z.; Wang, G.; Wang, Q.; Gao, F. Learning from Students: Online Contrastive Distillation Network for General Continual Learning.

19. Kumaran, D.; Hassabis, D.; McClelland, J.L. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in cognitive sciences* **2016**, *20*, 512–534.

20. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research* **2020**, *21*, 5485–5551.

21. Bellitto, G.; Pennisi, M.; Palazzo, S.; Bonicelli, L.; Boschini, M.; Calderara, S.; Spampinato, C. Effects of Auxiliary Knowledge on Continual Learning. *arXiv preprint arXiv:2206.02577* **2022**.

22. Xu, J.; Zhu, Z. Reinforced continual learning. *Advances in Neural Information Processing Systems* **2018**, *31*.

23. Marsocci, V.; Scardapane, S. Continual barlow twins: continual self-supervised learning for remote sensing semantic segmentation. *IEEE J-STARS* **2023**.

24. Derakhshani, M.M.; Zhen, X.; Shao, L.; Snoek, C. Kernel continual learning. In Proceedings of the International Conference on Machine Learning. PMLR, 2021, pp. 2621–2631.

25. Clem, R.L.; Celikel, T.; Barth, A.L. Ongoing in vivo experience triggers synaptic metaplasticity in the neocortex. *Science* **2008**, *319*, 101–104.

26. Zhai, M.; Chen, L.; Tung, F.; He, J.; Nawhal, M.; Mori, G. Lifelong gan: Continual learning for conditional image generation. In Proceedings of the Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 2759–2768.

27. Liu, X.; Hu, Y.S.; Cao, X.S.; Bagdanov, A.D.; Li, K.; Cheng, M.M. Long-Tailed Class Incremental Learning. In Proceedings of the European Conference on Computer Vision. Springer, 2022, pp. 495–512.

28. Khan, S.; Naseer, M.; Hayat, M.; Zamir, S.W.; Khan, F.S.; Shah, M. Transformers in vision: A survey. *ACM Computing Surveys* **2022**, *54*, 1–41.

29. Von Oswald, J.; Zhao, D.; Kobayashi, S.; Schug, S.; Caccia, M.; Zucchet, N.; Sacramento, J. Learning where to learn: Gradient sparsity in meta and continual learning. *Advances in Neural Information Processing Systems* **2021**, *34*, 5250–5263.

30. Cossu, A.; Tuytelaars, T.; Carta, A.; Passaro, L.; Lomonaco, V.; Bacciu, D. Continual Pre-Training Mitigates Forgetting in Language and Vision. *arXiv preprint arXiv:2205.09357* **2022**.

31. Borsos, Z.; Mutny, M.; Krause, A. Coresets via bilevel optimization for continual learning and streaming. *Advances in Neural Information Processing Systems* **2020**, *33*, 14879–14890.

32. Geishauser, C.; van Niekerk, C.; Lin, H.C.; Lubis, N.; Heck, M.; Feng, S.; Gasic, M. Dynamic Dialogue Policy for Continual Reinforcement Learning. In Proceedings of the Proceedings of the 29th International Conference on Computational Linguistics, 2022, pp. 266–284.

33. Doan, T.; Mirzadeh, S.I.; Pineau, J.; Farajtabar, M. Efficient Continual Learning Ensembles in Neural Network Subspaces. *arXiv preprint arXiv:2202.09826* **2022**.

34. Ye, F.; Bors, A.G. Task-Free Continual Learning via Online Discrepancy Distance Learning. *arXiv preprint arXiv:2210.06579* **2022**.

35. Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In Proceedings of the International Conference on Learning Representations, 2020.

36. Yan, S.; Xie, J.; He, X. Der: Dynamically expandable representation for class incremental learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 3014–3023.

37. Efrat, A.; Levy, O. The Turking Test: Can Language Models Understand Instructions? *arXiv preprint arXiv:2010.11982* **2020**.

38. Kulesza, A.; Taskar, B.; et al. Determinantal point processes for machine learning. *Foundations and Trends® in Machine Learning* **2012**, *5*, 123–286.

39. Doshi, K.; Yilmaz, Y. Continual learning for anomaly detection in surveillance videos. In Proceedings of the CVPR Workshops, 2020, pp. 254–255.

40. Radford, A.; Kim, J.W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. Learning transferable visual models from natural language supervision. In Proceedings of the International Conference on Machine Learning. PMLR, 2021, pp. 8748–8763.

41. Rostami, M.; Kolouri, S.; Pilly, P.; McClelland, J. Generative continual concept learning. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2020, Vol. 34, pp. 5545–5552.

42. Fernando, C.; Banarse, D.; Blundell, C.; Zwols, Y.; Ha, D.; Rusu, A.A.; Pritzel, A.; Wierstra, D. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734* **2017**.

43. Sun, J.; Wang, S.; Zhang, J.; Zong, C. Distill and replay for continual language learning. In Proceedings of the Proceedings of the 28th international conference on computational linguistics, 2020, pp. 3569–3579.

44. Rostami, M. Lifelong domain adaptation via consolidated internal distribution. *Advances in Neural Information Processing Systems* **2021**, *34*, 11172–11183.

45. Deng, D.; Chen, G.; Hao, J.; Wang, Q.; Heng, P.A. Flattening Sharpness for Dynamic Gradient Projection Memory Benefits Continual Learning. *Advances in Neural Information Processing Systems* **2021**, *34*, 18710–18721.

46. Ahn, H.; Cha, S.; Lee, D.; Moon, T. Uncertainty-based continual learning with adaptive regularization. *Advances in Neural Information Processing Systems* **2019**, *32*.

47. Caccia, L.; Belilovsky, E.; Caccia, M.; Pineau, J. Online learned continual compression with adaptive quantization modules. In Proceedings of the International Conference on Machine Learning. PMLR, 2020, pp. 1240–1250.

48. Zniyed, Y.; Nguyen, T.P.; et al. Efficient tensor decomposition-based filter pruning. *Neural Networks* **2024**, *178*, 106393.

49. Wang, L.; Zhang, X.; Yang, K.; Yu, L.; Li, C.; Hong, L.; Zhang, S.; Li, Z.; Zhong, Y.; Zhu, J. Memory Replay with Data Compression for Continual Learning. In Proceedings of the International Conference on Learning Representations, 2021.

50. Iscen, A.; Zhang, J.; Lazebnik, S.; Schmid, C. Memory-efficient incremental learning through feature adaptation. In Proceedings of the European Conference on Computer Vision. Springer, 2020, pp. 699–715.

51. Liu, M.; Chang, S.; Huang, L. Incremental Prompting: Episodic Memory Prompt for Lifelong Event Detection. *arXiv preprint arXiv:2204.07275* **2022**.

52. Li, X.; Wang, W.; Wu, L.; Chen, S.; Hu, X.; Li, J.; Tang, J.; Yang, J. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in Neural Information Processing Systems* **2020**, *33*, 21002–21012.

53. Doan, T.; Bennani, M.A.; Mazoure, B.; Rabusseau, G.; Alquier, P. A theoretical analysis of catastrophic forgetting through the ntk overlap matrix. In Proceedings of the International Conference on Artificial Intelligence and Statistics. PMLR, 2021, pp. 1072–1080.

54. Bennani, M.A.; Doan, T.; Sugiyama, M. Generalisation guarantees for continual learning with orthogonal gradient descent. *arXiv preprint arXiv:2006.11942* **2020**.

55. Gopalakrishnan, S.; Singh, P.R.; Fayek, H.; Ramasamy, S.; Ambikapathi, A. Knowledge capture and replay for continual learning. In Proceedings of the Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2022, pp. 10–18.

56. Lin, H.; Zhang, Y.; Qiu, Z.; Niu, S.; Gan, C.; Liu, Y.; Tan, M. Prototype-guided continual adaptation for class-incremental unsupervised domain adaptation. In Proceedings of the European Conference on Computer Vision. Springer, 2022, pp. 351–368.

57. Wu, C.; Herranz, L.; Liu, X.; van de Weijer, J.; Raducanu, B.; et al. Memory replay gans: Learning to generate new categories without forgetting. *Advances in Neural Information Processing Systems* **2018**, *31*.

58. Mirzadeh, S.I.; Farajtabar, M.; Gorur, D.; Pascanu, R.; Ghasemzadeh, H. Linear Mode Connectivity in Multi-task and Continual Learning. In Proceedings of the International Conference on Learning Representations, 2020.

59. Zhou, M.; Xiao, J.; Chang, Y.; Fu, X.; Liu, A.; Pan, J.; Zha, Z.J. Image de-raining via continual learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 4907–4916.

60. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* **2021**.

61. Wang, L.; Zhang, M.; Jia, Z.; Li, Q.; Bao, C.; Ma, K.; Zhu, J.; Zhong, Y. Afec: Active forgetting of negative transfer in continual learning. *Advances in Neural Information Processing Systems* **2021**, *34*, 22379–22391.

62. Yang, G.; Pan, F.; Gan, W.B. Stably maintained dendritic spines are associated with lifelong memories. *Nature* **2009**, *462*, 920–924.

63. Biesialska, M.; Biesialska, K.; Costa-jussà, M.R. Continual Lifelong Learning in Natural Language Processing: A Survey. In Proceedings of the Proceedings of the 28th International Conference on Computational Linguistics, 2020, pp. 6523–6541.

64. Peng, B.; Risteski, A. Continual learning: a feature extraction formalization, an efficient algorithm, and fundamental obstructions. *arXiv preprint arXiv:2203.14383* **2022**.

65. Zhao, J.; Zhang, X.; Zhao, B.; Hu, W.; Diao, T.; Wang, L.; Zhong, Y.; Li, Q. Genetic dissection of mutual interference between two consecutive learning tasks in Drosophila. *Elife* **2023**, *12*, e83516.

66. Singh, P.; Mazumder, P.; Rai, P.; Namboodiri, V.P. Rectification-based knowledge retention for continual learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 15282–15291.

67. Carr, M.F.; Jadhav, S.P.; Frank, L.M. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature neuroscience* **2011**, *14*, 147–153.

68. Douillard, A.; Cord, M.; Ollion, C.; Robert, T.; Valle, E. Podnet: Pooled outputs distillation for small-tasks incremental learning. In Proceedings of the European Conference on Computer Vision. Springer, 2020, pp. 86–102.

69. Yan, S.; Hong, L.; Xu, H.; Han, J.; Tuytelaars, T.; Li, Z.; He, X. Generative Negative Text Replay for Continual Vision-Language Pretraining. In Proceedings of the European Conference on Computer Vision. Springer, 2022, pp. 22–38.

70. Ye, F.; Bors, A.G. Learning latent representations across multiple data domains using lifelong vaegan. In Proceedings of the European Conference on Computer Vision. Springer, 2020, pp. 777–795.

71. Mi, F.; Kong, L.; Lin, T.; Yu, K.; Faltings, B. Generalized class incremental learning. In Proceedings of the CVPR Workshops, 2020, pp. 240–241.

72. Liu, T.; Ungar, L.; Sedoc, J. Continual Learning for Sentence Representations Using Conceptors. In Proceedings of the Proceedings of NAACL-HLT, 2019, pp. 3274–3279.

73. Hayashi-Takagi, A.; Yagishita, S.; Nakamura, M.; Shirai, F.; Wu, Y.I.; Loshbaugh, A.L.; Kuhlman, B.; Hahn, K.M.; Kasai, H. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* **2015**, *525*, 333–338.

74. Mirzadeh, S.I.; Chaudhry, A.; Yin, D.; Hu, H.; Pascanu, R.; Gorur, D.; Farajtabar, M. Wide neural networks forget less catastrophically. In Proceedings of the International Conference on Machine Learning. PMLR, 2022, pp. 15699–15717.

75. Wang, Y.; Huang, Z.; Hong, X. S-Prompts Learning with Pre-trained Transformers: An Occam's Razor for Domain Incremental Learning. *arXiv preprint arXiv:2207.12819* **2022**.

76. Hou, S.; Pan, X.; Loy, C.C.; Wang, Z.; Lin, D. Learning a unified classifier incrementally via rebalancing. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 831–839.

77. Petit, G.; Popescu, A.; Schindler, H.; Picard, D.; Delezoide, B. FeTrIL: Feature Translation for Exemplar-Free Class-Incremental Learning. *arXiv preprint arXiv:2211.13131* **2022**.

78. Zhang, C.; Song, N.; Lin, G.; Zheng, Y.; Pan, P.; Xu, Y. Few-shot incremental learning with continually evolved classifiers. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 12455–12464.

79. Dong, S.; Hong, X.; Tao, X.; Chang, X.; Wei, X.; Gong, Y. Few-shot class-incremental learning via relation knowledge distillation. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2021, Vol. 35, pp. 1255–1263.

80. Izmailov, P.; Podoprikhin, D.; Garipov, T.; Vetrov, D.; Wilson, A.G. Averaging weights leads to wider optima and better generalization. *arXiv preprint arXiv:1803.05407* **2018**.

81. Houlsby, N.; Giurgiu, A.; Jastrzebski, S.; Morrone, B.; De Laroussilhe, Q.; Gesmundo, A.; Attariyan, M.; Gelly, S. Parameter-efficient transfer learning for NLP. In Proceedings of the International Conference on Machine Learning. PMLR, 2019, pp. 2790–2799.

82. Mazumder, P.; Singh, P.; Rai, P. Few-shot lifelong learning. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2021, Vol. 35, pp. 2337–2345.

83. Liu, H.; Gu, L.; Chi, Z.; Wang, Y.; Yu, Y.; Chen, J.; Tang, J. Few-Shot Class-Incremental Learning via Entropy-Regularized Data-Free Replay. In Proceedings of the European Conference on Computer Vision. Springer, 2022, pp. 146–162.

84. Li, A.; Boyd, A.; Smyth, P.; Mandt, S. Variational Beam Search for Learning with Distribution Shifts. *arXiv preprint arXiv:2012.08101* **2020**.

85. Mirzadeh, S.I.; Farajtabar, M.; Pascanu, R.; Ghasemzadeh, H. Understanding the role of training regimes in continual learning. *Advances in Neural Information Processing Systems* **2020**, *33*, 7308–7320.

86. Jha, S.; Schiemer, M.; Ye, J. Continual learning in human activity recognition: an empirical analysis of regularization. *arXiv preprint arXiv:2007.03032* **2020**.

87. Banayeeanzade, M.; Mirzaiezadeh, R.; Hasani, H.; Soleymani, M. Generative vs. Discriminative: Rethinking The Meta-Continual Learning. *Advances in Neural Information Processing Systems* **2021**, *34*, 21592–21604.

88. Berga, D.; Masana, M.; Van de Weijer, J. Disentanglement of color and shape representations for continual learning. *arXiv preprint arXiv:2007.06356* **2020**.

89. Saha, G.; Garg, I.; Roy, K. Gradient Projection Memory for Continual Learning. In Proceedings of the International Conference on Learning Representations, 2020.

90. Qin, C.; Joty, S. Continual Few-shot Relation Learning via Embedding Space Regularization and Data Augmentation. In Proceedings of the Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2022, pp. 2776–2789.

91. Pham, V.T.; Zniyed, Y.; Nguyen, T.P. Enhanced network compression through tensor decompositions and pruning. *IEEE Transactions on Neural Networks and Learning Systems* **2024**.

92. Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; Tuytelaars, T. Memory aware synapses: Learning what (not) to forget. In Proceedings of the Proceedings of the European Conference on Computer Vision, 2018, pp. 139–154.

93. Wołczyk, M.; Piczak, K.; Wójcik, B.; Pustelnik, L.; Morawiecki, P.; Tabor, J.; Trzcinski, T.; Spurek, P. Continual Learning with Guarantees via Weight Interval Constraints. In Proceedings of the International Conference on Machine Learning. PMLR, 2022, pp. 23897–23911.

94. Ramasesh, V.V.; Dyer, E.; Raghu, M. Anatomy of catastrophic forgetting: Hidden representations and task semantics. *arXiv preprint arXiv:2007.07400* **2020**.

95. Fini, E.; da Costa, V.G.T.; Alameda-Pineda, X.; Ricci, E.; Alahari, K.; Mairal, J. Self-supervised models are continual learners. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 9621–9630.

96. Isele, D.; Cosgun, A. Selective experience replay for lifelong learning. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2018, Vol. 32.

97. González, O.C.; Sokolov, Y.; Krishnan, G.P.; Delanois, J.E.; Bazhenov, M. Can sleep protect memories from catastrophic forgetting? *Elife* **2020**, *9*.

98. Pasunuru, R.; Stoyanov, V.; Bansal, M. Continual Few-Shot Learning for Text Classification. In Proceedings of the Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 5688–5702.

99. Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A.A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **2017**, *114*, 3521–3526.

100. Kang, Z.; Fini, E.; Nabi, M.; Ricci, E.; Alahari, K. A soft nearest-neighbor framework for continual semi-supervised learning. *arXiv preprint arXiv:2212.05102* **2022**.

101. Liu, Q.; Yu, X.; He, S.; Liu, K.; Zhao, J. Lifelong Intent Detection via Multi-Strategy Rebalancing. *arXiv preprint arXiv:2108.04445* **2021**.

102. Wang, R.; Bao, Y.; Zhang, B.; Liu, J.; Zhu, W.; Guo, G. Anti-retroactive interference for lifelong learning. In Proceedings of the European Conference on Computer Vision. Springer, 2022, pp. 163–178.

103. Stan, S.; Rostami, M. Unsupervised model adaptation for continual semantic segmentation. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2021, Vol. 35, pp. 2593–2601.

104. Hocquet, G.; Bichler, O.; Querlioz, D. Ova-inn: Continual learning with invertible neural networks. In Proceedings of the International Joint Conference on Neural Networks. IEEE, 2020, pp. 1–7.