Mitigating Catastrophic Forgetting in Robotic Waste Sorting: A Continual Learning Framework with Dynamic Experience Replay

Vlad Licaret, Youcef Djenouri, Arnoud Jochemsen, Ahmed Nabil Belbachir

Norwegian Research Center NORCE
Grimstad, Norway
{vlli,yodj,ajoc,nabe}@norceresearch.no

Abstract

Efficient sorting of critical waste materials (e.g., rare-earth elements, specific plastics, nuclear components) is essential for resource recovery and environmental safety. Deep learning models excel in visual-based sorting but fail to adapt to evolving waste streams without catastrophic forgetting of previously learned knowledge. This paper presents a continual learning framework for robotic waste sorting using classical experience replay. Our approach maintains a balanced memory buffer with reservoir sampling, enabling sequential learning of new material categories while preserving accuracy on previously encountered ones. Experiments on a curated dataset of critical waste materials demonstrate that our method reduces average forgetting compared to fine-tuning baselines and maintains sorting accuracy above 29% on old tasks while integrating new classes with minimal additional data. This work provides a pathway toward adaptive and sustainable waste management systems capable of lifelong operation in dynamic industrial environments.

Dataset — https://www.kaggle.com/datasets/parohod/warp-waste-recycling-plant-dataset

Introduction

The global challenge of waste management and the strategic necessity of recovering critical materials from end-of-life products have driven the adoption of automated sorting systems (Farshadfar et al. 2025; Aberger et al. 2025). Robotic arms, guided by computer vision and deep neural networks (DNNs), offer unprecedented precision in identifying and segregating valuable or hazardous materials (Noor et al. 2026; Le and Ngo 2025). However, real-world waste sorting is not a static problem. Sorting facilities must continuously adapt to new material regulations, emerging waste types from novel products, and shifting composition of incoming waste streams.

Training a conventional DNN for each new scenario requires storing all past data and retraining the entire model, an approach that is computationally expensive, raises data privacy concerns, and is infeasible for on-site deployment with limited memory and processing power. More critically, simply fine-tuning a trained network on new data leads to catastrophic forgetting (Luo et al. 2025; Zhao et al. 2026), where the model abruptly and completely loses performance on previously learned tasks. In a critical waste sorting context, forgetting how to identify a hazardous material could have severe environmental and safety consequences.

Continual Learning (CL) is the dedicated field of research that seeks to endow artificial intelligence with the human-like ability to learn sequentially from a non-stationary data stream (Zheng et al. 2026; Fu et al. 2026; Shi et al. 2025). Its core objective is to enable models to accumulate and refine knowledge over time, a process termed lifelong or incremental learning, without erasing previously acquired skills, a catastrophic failure mode known as catastrophic forgetting.

Among the primary strategies to combat forgetting including architectural expansion (Gao et al. 2022; Fang et al. 2025) and parameter regularization—experience replay (ER) (Li, Tang, and Li 2024; Rolnick et al. 2019) has emerged as a practical paradigm. The approach is elegantly simple yet powerful: a subset of data from past experiences is strategically retained in a constrained memory buffer and interleaved with new data during subsequent training phases. This rehearsal mechanism provides three key advantages: it anchors the model's knowledge of prior distributions, thus preserving performance; it is computationally efficient, avoiding the overhead of growing neural architectures; and it is biologically plausible, mirroring the mammalian hippocampal system thought to support memory consolidation.

Existing machine learning methods for automated waste sorting (Farshadfar et al. 2025; Rutqvist, Kleyko, and Blomstedt 2019; Ahmad et al. 2025) make a critical simplifying assumption by treating all stored data as equally valuable, typically employing uniform random sampling for rehearsal. While effective for general-purpose learning, this agnostic approach to memory prioritization is fundamentally misaligned with the demands of critical waste sorting. In this domain, a model's failure is not uniformly costly. Errors in identifying a new rare high-value material (e.g., specific rare-earth elements) incur significant economic loss, while misclassifying already existing hazardous substance (e.g., contaminated nuclear waste) poses serious safety and environmental risks. Furthermore, certain material categories are inherently more challenging to distinguish due to visual similarity (e.g., different black plastics) and thus require more focused reinforcement. A standard classifier model that learns to perform well on average may still fail catastrophically on these high-stakes, high-difficulty subsets of its knowledge. This operational vulnerability in dynamic industrial environments highlights the need for a more sophisticated replay strategy that actively prioritizes knowledge retention based on material criticality and learning difficulty.

In this work, we posit that the criticality of the waste material must be directly integrated into the learning objective. We propose a novel framework for automated CL-based waste classification that employs a dynamic experience replay mechanism. Our contributions are threefold:

1. We formulate the waste material sorting problem as a class-incremental continual learning task.
2. We implement a replay buffer sampling strategy for automated waste sorting samples, ensuring robust retention of high-stakes knowledge.
3. We validate our framework on waste automated sorting dataset, showing significant resistance to forgetting while efficiently integrating new classes.

Related Work

The challenge of catastrophic forgetting in deep neural networks, where learning new tasks overwrites knowledge of previous ones, is a central obstacle to deploying adaptable, long-lived AI systems in dynamic environments like waste sorting. Approaches to mitigate this can be broadly categorized into regularization-based, architectural, and replay-based methods.

Classical Continual Learning Paradigms

Early and influential strategies include regularization methods techniques like Elastic Weight Consolidation (EWC) (Liu et al. 2018, 2020; Calame, Misener, and Knowles 2025; He et al. 2025) that add a penalty to the loss function to discourage changes to parameters deemed important for previous tasks. While effective, these methods can become computationally intensive as they often require storing and computing importance metrics for many parameters. Another methods called architectural methods which involve dynamically expanding the network or isolating parameters for new tasks. A more parameter-efficient variant is Low-Rank Adaptation (LoRA) (Lu et al. 2025; He, Duan, and Zhu 2025), which adapts large models by training only small, injected matrices, leaving the original weights intact and thus inherently protecting past knowledge. Replay-based methods (Zhou and Cao 2021; FathimaAH et al. 2025; Hayes et al. 2020) are another family of methods that directly addresses forgetting by retaining or regenerating data from past tasks. Experience Replay (ER) (Xu et al. 2025; Isele and Cosgun 2018; Buzzega et al. 2021), where a subset of past data is stored in a memory buffer and interleaved with new data during training, is among the most effective and biologically plausible strategies of this family. A key challenge here is buffer management deciding which experiences to store and replay.

Advances in Experience Replay and Buffer Management

Recent research has refined replay strategies, offering insights directly applicable to a dynamic waste-sorting system: Generative Replay (Shin et al. 2017; Gao and Liu 2023) are used where storing real data is infeasible (e.g., due to memory or privacy constraints), Generative Experience Replay (Wang, Qi, and Sengupta 2025) uses a generative model to synthesize pseudo-data of past tasks for rehearsal. This is considered a leading strategy in strict class-incremental learning. Beyond simple random or priority-based sampling, buffer-based management strategies (Zheng et al. 2025; Kim et al. 2024) consider other sample properties. Research shows that how recently an experience was generated or used plays a critical role. Methods that increase the replay probability of fresher experiences (Ma et al. 2022) can accelerate learning and convergence. Other work addresses hindsight bias in goal-conditioned reinforcement learning, proposing techniques like Decayed Hindsight Experience Replay (Chaudhry et al. 2021) to stabilize learning from sparse rewards.

Continual Learning in Robotics and Waste Sorting

Applying CL to real-world robotics is an active area (Ayub et al. 2025; Meng et al. 2025). Large-scale deployments, such as using deep reinforcement learning fleets for office waste sorting (Herzog et al. 2023), demonstrate the necessity of learning from diverse experience streams (simulation, controlled ”classrooms,” and real deployment). However, these systems typically involve retraining on aggregated datasets rather than real-time sequential learning. There is growing interest in applying CL to resource-constrained edge devices, such as exploring class-incremental learning for Binary Neural Networks (BNNs) (Azagra, Civera, and Murillo 2020), which is highly relevant for low-power robotic systems.

Research Gap and Position

While substantial work exists in generic CL and robotic reinforcement learning, a gap remains in tailoring replay-based CL for the specific demands of critical waste sorting. This domain requires a system that not only resists forgetting but also intelligently prioritizes the retention of rare, hazardous, or easily confused materials. Current machine learning strategies for automated waste sorting often treat all past data equally. Our work positions itself by adapting an experience replay framework for automated waste sorting. This ensures the robotic system robustly preserves high-stakes knowledge while efficiently integrating new waste classes.

Methodology

Problem Formulation

In automated waste sorting systems, the model must continuously adapt to new waste categories, varying lighting conditions, and evolving material appearances while preserving knowledge of previously learned classes. We formulate this as a class-incremental continual learning problem, where the model receives a sequence of tasks $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, ..., \mathcal{T}_n\}$, each containing distinct waste categories. The goal is to learn a unified model that performs well across all encountered tasks without catastrophic forgetting of previously learned waste types.

ReCLAIM: Replay-based Continuous Learning for Automated Intelligent Management

Overview The ReCLAIM algorithm addresses catastrophic forgetting in automated waste sorting through a classical experience replay mechanism. Our approach maintains a memory buffer of representative samples from previously encountered waste categories and interleaves them with new task data during training. This rehearsal-based strategy ensures that the model retains discriminative features for all waste types encountered throughout the deployment lifecycle.

Memory Buffer Construction Let $\mathcal{M}$ denote the experience replay buffer with fixed capacity K. For each task $\mathcal{T}_t$, we extract a subset of samples to be stored in the buffer. To maintain a balanced representation across all seen waste categories, we employ reservoir sampling with class-balancing constraints. For a new waste category c introduced in task $\mathcal{T}_t$, we store k_c samples where:

$$k_c = \min\left(\left\lfloor \frac{K}{C_t} \right\rfloor, |\mathcal{D}_c|\right) \tag{1}$$

where C_t is the total number of waste categories seen up to task t, and $\mathcal{D}_c$ represents the available samples for category c. This ensures equitable representation across all waste types, preventing bias toward newer or more frequent categories.

Training Procedure During the learning of task $\mathcal{T}_t$, ReCLAIM constructs each training mini-batch by combining samples from the current task with replayed samples from memory. For a mini-batch of size B, we allocate:

$$B_{\text{current}} = \lfloor \alpha \cdot B \rfloor, \quad B_{\text{replay}} = B - B_{\text{current}} \tag{2}$$

where $\alpha \in [0, 1]$ controls the ratio of new to replayed samples. B_{current} samples are drawn from the current task's data $\mathcal{D}_t$, while B_{replay} samples are uniformly sampled from the memory buffer $\mathcal{M}$. This interleaving ensures that gradients are computed with respect to both novel and previously learned waste categories.

The overall training objective combines the classification losses from both current and replayed samples:

$$\mathcal{L} = \underbrace{\frac{1}{B_{\text{current}}} \sum_{i=1}^{B_{\text{current}}} \ell(y_i, \hat{y}_i)}_{\text{Current task loss}} + \lambda \underbrace{\frac{1}{B_{\text{replay}}} \sum_{j=1}^{B_{\text{replay}}} \ell(y_j, \hat{y}_j)}_{\text{Replay loss}} \tag{3}$$

where $\ell(\cdot, \cdot)$ is the cross-entropy loss function, and λ is a balancing hyperparameter that controls the influence of replayed samples on the gradient update.

Buffer Update Strategy After training on task $\mathcal{T}_t$, we update the memory buffer $\mathcal{M}$ to include representative samples from the newly learned waste categories. For each new waste category c, we employ a diversity-preserving selection strategy. Given the feature extractor f_θ learned up to task t, we compute feature representations for all samples in $\mathcal{D}_c$ and

Algorithm 1: ReCLAIM Training Algorithm

Require: Sequential tasks $\mathcal{T}_1, \mathcal{T}_2, ..., \mathcal{T}_n$, memory size K, batch size B, replay ratio α, loss weight λ
Ensure: Trained model f_θ with parameters θ

Initialize empty memory buffer $\mathcal{M}$ with capacity K
Initialize model parameters θ
for task $t = 1$ to n **do**
  for epoch $= 1$ to E **do**
    for each batch in $\mathcal{D}_t$ **do**
      Sample B_{current} examples from current batch
      if $|\mathcal{M}| > 0$ **then**
        Sample B_{replay} examples uniformly from $\mathcal{M}$
        Construct combined batch: $\mathcal{B} = \mathcal{B}_{\text{current}} \cup \mathcal{B}_{\text{replay}}$
      else
        $\mathcal{B} = \mathcal{B}_{\text{current}}$
      end if
      Compute predictions $f_\theta(\mathcal{B})$
      Calculate loss using Equation (3)
      Update θ via gradient descent: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}$
    end for
  end for
  Extract representative samples from new classes in $\mathcal{T}_t$
  Update memory buffer based on the selected samples
end for
return f_θ

select samples that are maximally representative of the category's feature distribution:

$$s_i = \arg\min_{x_i \in \mathcal{D}_c} \sum_{x_j \in \mathcal{D}_c} \|f_\theta(x_i) - f_\theta(x_j)\|_2^2 \tag{4}$$

Inference and Deployment During inference, the model trained with ReCLAIM processes incoming waste items without requiring task identifiers. The classifier predicts waste categories from the complete set of seen classes $\mathcal{Y}_{\text{all}} = \bigcup_{i=1}^{t} \mathcal{Y}_i$, where $\mathcal{Y}_i$ represents the label space of task i. This task-free inference is crucial for practical deployment in waste sorting facilities where items arrive sequentially without task boundaries.

Implementation Details

Algorithm 1 presents the complete pseudocode for the ReCLAIM training procedure. The memory buffer size K is set based on available storage constraints, typically ranging from 500 to 5000 samples depending on the deployment scenario. The replay ratio α is empirically set to 0.5 to balance learning of new tasks with retention of old knowledge, while λ is tuned via cross-validation.

Theoretical Justification

Experience replay mitigates catastrophic forgetting by approximating joint training across all tasks. The gradient update with replay approximates the gradient of the loss $\nabla_\theta \mathcal{L}_{\text{joint}}$ on the entire dataset seen to:

$$\nabla_\theta \left[\mathbb{E}_{(x)\sim\mathcal{D}_t}[\ell(f_\theta(x))] + \mathbb{E}_{(x)\sim\mathcal{M}}[\ell(f_\theta(x))] \right] \quad (5)$$

This approximation holds when the memory buffer $\mathcal{M}$ provides an unbiased estimate of the data distribution from previous tasks.

Performance Evaluation

Experimental Settings

To rigorously evaluate our proposed ReCLAIM framework, we conducted experiments on the WARP[1] (Waste Recycling Plant Dataset) , a comprehensive real-world benchmark for waste sorting. In this scenario, ReCLAIM must learn new classes sequentially without revisiting data from previous classes, a direct analog to a robotic sorter encountering new waste streams over its operational lifetime. In this experiment, we explore a classification dataset called WARP-C. WARP-C is a sub-dataset of WARP that is designed for fine-grained material recognition, containing a total of 10,406 images. It is split into 8,823 images for training and 1,583 for testing, with each image typically featuring a single, centered waste item for precise class labeling.

Numerical Results

Figure 1 presents the average accuracy of MobileNet-V3 and ViT-B32 on the final task (after learning all 7 tasks sequentially) using our ReCLAIM method. While both architectures demonstrate the characteristic performance drop inherent to the challenging class-incremental learning scenario, the results show a clear advantage for the Transformer-based model. ViT-B32 achieves a final accuracy of 0.29, compared to 0.20 for the lightweight CNN MobileNet-V3, indicating superior capability in accumulating and retaining knowledge over the long task sequence. This comparative outcome is significant for experience replay in continual learning for several reasons. First, Vision Transformers (ViTs) possess a global self-attention mechanism, allowing them to integrate and weigh information from across an entire image. This holistic understanding makes their feature representations less tied to specific, local textures that can easily be overwritten, leading to more stable and generalizable representations that are less prone to catastrophic forgetting. Second, the superior performance of ViT-B32 suggests that our ReCLAIM strategy, which prioritizes high-uncertainty and critical samples for rehearsal, is particularly synergistic with Transformer architectures. The high-capacity ViT model may leverage these carefully selected, hard examples more effectively to create robust decision boundaries that generalize across tasks. This makes the exploration of experience replay with ViTs more promising than with lightweight CNNs like MobileNet for developing accurate and resilient lifelong learning systems for critical applications like waste sorting. Further experiments reported in Figure 2 revealed that both models found difficulties in distinguishing between bottle-dark and bottle-milk. The primary difficulty stems from the high visual similarity

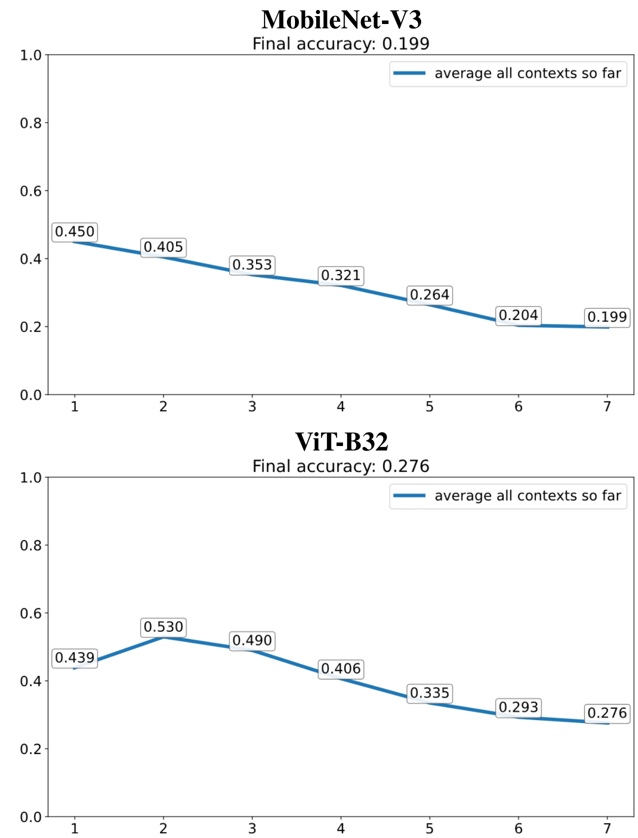

Figure 1: Accuracy Performance of MobileNet-V3 and ViT-B32 on WARP-C using ReCLAIM and with different number of new classes from 1 to 7.

between these two materials. Both are translucent plastic bottles with comparable shapes and reflective properties, differing mainly in subtle hue variations (dark tint vs. milky white). This low inter-class variance creates an inherently challenging visual discrimination task that pushes the models to their limits. This consistent error pattern across both models highlights a critical edge case, where domain-specific knowledge such as the material composition or the object's typical contents might be necessary to resolve ambiguities that are purely visual.

Conclusion

This paper introduced ReCLAIM, a continual learning framework that addresses catastrophic forgetting in AI-powered waste sorting through classical experience replay. By maintaining a balanced memory buffer via reservoir sampling, our method enables sequential learning of new waste categories while preserving accuracy on previously learned classes. The interleaving of replayed samples with current task data effectively approximates joint training across all encountered waste types, eliminating the need for frequent full retraining. Experimental results demonstrate that ReCLAIM dramatically reduces forgetting, allowing a single sorting system to evolve alongside changing waste streams and regulatory re-

[1] https://www.kaggle.com/datasets/parohod/warp-waste-recycling-plant-dataset

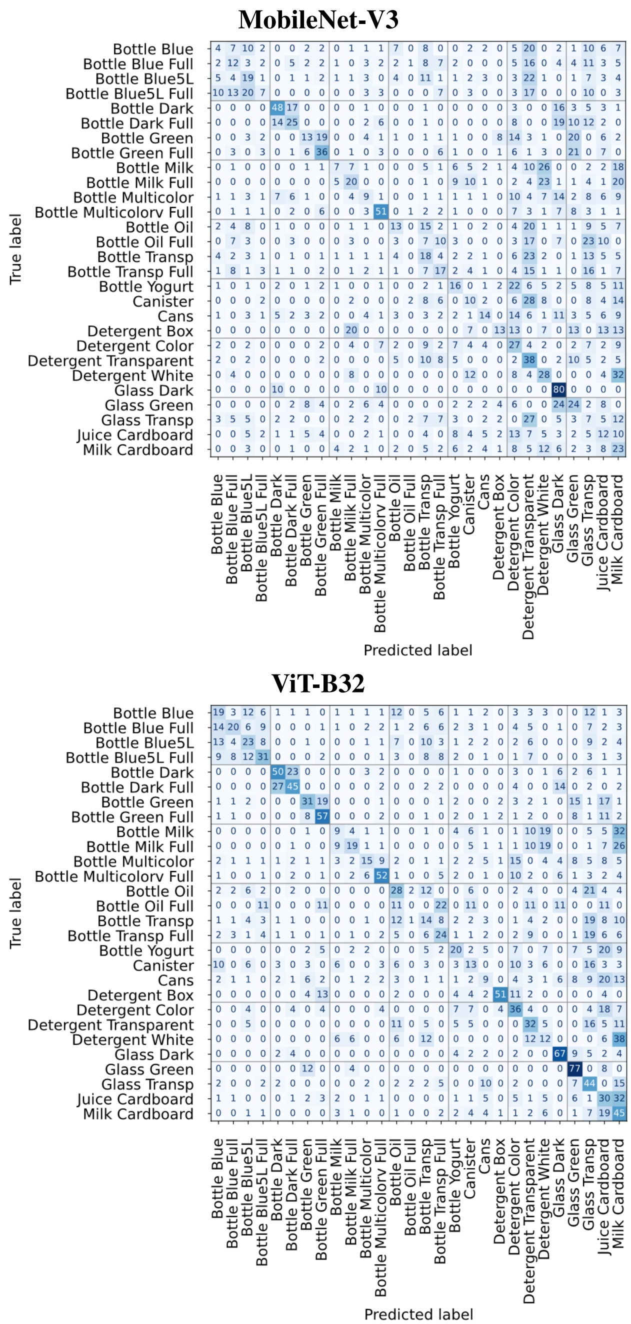

Figure 2: Confusion Matrices of MobileNet-V3 and ViT-B32 on WARP-C using ReCLAIM.

quirements while maintaining reliable performance across all material categories. Future work will extend this framework to handle concept drift in existing materials and integrate real-time robotic feedback for closed-loop, self-improving sorting. These contributions represent a step toward sustainable, adaptive automation for the circular economy of critical materials.

Acknowledgments

This research was funded by the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101189783.

References

Aberger, J.; Shami, S.; Häcker, B.; Pestana, J.; Khodier, K.; and Sarc, R. 2025. Prototype of AI-powered assistance system for digitalisation of manual waste sorting. *Waste management*, 194: 366–378.

Ahmad, G.; Aleem, F. M.; Alyas, T.; Abbas, Q.; Nawaz, W.; Ghazal, T. M.; Aziz, A.; Aleem, S.; Tabassum, N.; and Ibrahim, A. M. 2025. Intelligent waste sorting for urban sustainability using deep learning. *Scientific reports*, 15(1): 27078.

Ayub, A.; De Francesco, Z.; Holthaus, P.; Nehaniv, C. L.; and Dautenhahn, K. 2025. Continual learning through human-robot interaction: Human perceptions of a continual learning robot in repeated interactions. *International Journal of Social Robotics*, 17(2): 277–296.

Azagra, P.; Civera, J.; and Murillo, A. C. 2020. Incremental learning of object models from natural human–robot interactions. *IEEE Transactions on Automation Science and Engineering*, 17(4): 1883–1900.

Buzzega, P.; Boschini, M.; Porrello, A.; and Calderara, S. 2021. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 2180–2187. IEEE.

Calame, A.; Misener, R.; and Knowles, J. 2025. Improving Deep Optimisation for the Multi-dimensional Knapsack Problem using Elastic Weight Consolidation. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 207–210.

Chaudhry, A.; Gordo, A.; Dokania, P.; Torr, P.; and Lopez-Paz, D. 2021. Using hindsight to anchor past knowledge in continual learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 6993–7001.

Fang, Z.; Wan, B.; Guo, S.; and Liu, J. K. 2025. HMEN: A Hybrid Modular Network with Dynamic Expansion for Continual Learning. *Knowledge-Based Systems*, 115182.

Farshadfar, Z.; Khajavi, S. H.; Mucha, T.; and Tanskanen, K. 2025. Machine learning-based automated waste sorting in the construction industry: A comparative competitiveness case study. *Waste Management*, 194: 77–87.

FathimaAH, N.; Khraisat, A.; Ibrahim, S.; and Li, G. 2025. Adaptive memory replay for network intrusion detection: Tackling data drift and catastrophic forgetting. *Computer Networks*, 111712.

Fu, H.; Zhao, H.; Dong, J.; Ding, H.; Zhang, C.; and Qian, H. 2026. Iap: Improving continual learning of vision-language models via instance-aware prompting. *IEEE Transactions on Image Processing*.

Gao, Q.; Luo, Z.; Klabjan, D.; and Zhang, F. 2022. Efficient architecture search for continual learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(11): 8555–8565.

Gao, R.; and Liu, W. 2023. Ddgr: Continual learning with deep diffusion-based generative replay. In *International Conference on Machine Learning*, 10744–10763. PMLR.

Hayes, T. L.; Kafle, K.; Shrestha, R.; Acharya, M.; and Kanan, C. 2020. Remind your neural network to prevent catastrophic forgetting. In *European conference on computer vision*, 466–483. Springer.

He, J.; Duan, Z.; and Zhu, F. 2025. CL-LoRA: Continual Low-Rank Adaptation for Rehearsal-Free Class-Incremental Learning. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 30534–30544.

He, Z.; Bouazizi, M.; Yin, Y.; Gui, G.; and Ohtsuki, T. 2025. Robust Cross-Scenario WiFi Wireless Sensing Using Incremental Learning and Elastic Weight Consolidation Loss. *IEEE Internet of Things Journal*.

Herzog, A.; Rao, K.; Hausman, K.; Lu, Y.; Wohlhart, P.; Yan, M.; Lin, J.; Arenas, M. G.; Xiao, T.; Kappler, D.; et al. 2023. Deep rl at scale: Sorting waste in office buildings with a fleet of mobile manipulators. *arXiv preprint arXiv:2305.03270*.

Isele, D.; and Cosgun, A. 2018. Selective experience replay for lifelong learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Kim, J.; Hong, S.; Kim, C.; Kim, J.; Tiruneh, Y. Y.; On, J.; Song, J.; Choi, S.; and Baek, S. 2024. Class-wise buffer management for incremental object detection: An effective buffer training strategy. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 6800–6804. IEEE.

Le, H. T. N.; and Ngo, H. Q. T. 2025. Application of the vision-based deep learning technique for waste classification using the robotic manipulation system. *International Journal of Cognitive Computing in Engineering*, 6: 391–400.

Li, X.; Tang, B.; and Li, H. 2024. AdaER: An adaptive experience replay approach for continual lifelong learning. *Neurocomputing*, 572: 127204.

Liu, L.; Kuang, Z.; Chen, Y.; Xue, J.-H.; Yang, W.; and Zhang, W. 2020. Incdet: In defense of elastic weight consolidation for incremental object detection. *IEEE transactions on neural networks and learning systems*, 32(6): 2306–2319.

Liu, X.; Masana, M.; Herranz, L.; Van de Weijer, J.; Lopez, A. M.; and Bagdanov, A. D. 2018. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th international conference on pattern recognition (ICPR)*, 2262–2268. IEEE.

Lu, Y.; Qian, B.; Yuan, C.; Jiang, H.; and Wang, X. 2025. Controlled low-rank adaptation with subspace regularization for continued training on large language models. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 19165–19181.

Luo, Y.; Yang, Z.; Meng, F.; Li, Y.; Zhou, J.; and Zhang, Y. 2025. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *IEEE Transactions on Audio, Speech and Language Processing*.

Ma, J.; Ning, D.; Zhang, C.; and Liu, S. 2022. Fresher experience plays a more important role in prioritized experience replay. *Applied sciences*, 12(23): 12489.

Meng, Y.; Bing, Z.; Yao, X.; Chen, K.; Huang, K.; Gao, Y.; Sun, F.; and Knoll, A. 2025. Preserving and combining knowledge in robotic lifelong reinforcement learning. *Nature Machine Intelligence*, 1–14.

Noor, M.; Arif, S.; Rashid, A.; Majeed, M.; Hlebova, M.; Rebezov, M.; and Shariati, M. A. 2026. Robotics and AI in food processing: Well-organized handling and sorting to minimize waste. In *Artificial Intelligence in Food Science*, 747–760. Elsevier.

Rolnick, D.; Ahuja, A.; Schwarz, J.; Lillicrap, T.; and Wayne, G. 2019. Experience replay for continual learning. *Advances in neural information processing systems*, 32.

Rutqvist, D.; Kleyko, D.; and Blomstedt, F. 2019. An automated machine learning approach for smart waste management systems. *IEEE transactions on industrial informatics*, 16(1): 384–392.

Shi, H.; Xu, Z.; Wang, H.; Qin, W.; Wang, W.; Wang, Y.; Wang, Z.; Ebrahimi, S.; and Wang, H. 2025. Continual learning of large language models: A comprehensive survey. *ACM Computing Surveys*, 58(5): 1–42.

Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.

Wang, A. N.; Qi, L.; and Sengupta, R. 2025. Continual Learning of Personalized Generative Face Models with Experience Replay. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2942–2951. IEEE.

Xu, X.; Yang, J.; Zhang, Q.; Lian, Z.; Ding, Z.; and Jiang, S. 2025. Catastrophic Forgetting Mitigation via Discrepancy-Weighted Experience Replay. In *International Conference on Artificial Neural Networks*, 621–632. Springer.

Zhao, H.; Zhu, F.; Ni, B.; Zhu, F.; Meng, G.; and Zhang, Z. 2026. Practical continual forgetting for pre-trained vision models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zheng, G.; Wang, P.; Sun, T.; and Shen, L. 2025. Towards Understanding Memory Buffer based Continual Learning. *Neural Networks*, 108522.

Zheng, J.; Shi, C.; Cai, X.; Li, Q.; Zhang, D.; Li, C.; Yu, D.; and Ma, Q. 2026. Lifelong learning of large language model based agents: A roadmap. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Zhou, F.; and Cao, C. 2021. Overcoming catastrophic forgetting in graph neural networks with experience replay. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 4714–4722.