CGR: Confidence-Guided Replay for Buffer-Based Continual Learning

Anonymous authors
Paper under double-blind review

Abstract

Continual Learning (CL) aims to acquire new knowledge while preserving previously learned information without catastrophic forgetting. Buffer-based methods, which retain samples from past tasks, have demonstrated promising results; however, efficiently allocating limited buffer space remains a significant challenge. Recent studies often either neglect the varying impact individual samples have on the learning process or incur high computational costs to identify informative replay samples. To overcome these limitations, we propose a novel approach called Confidence-Guided Replay (CGR), a lightweight policy that dynamically allocates the buffer by monitoring confidence fluctuations in the main continual learner model. Leveraging measures of sample contribution and difficulty, CGR adaptively prioritizes highly informative samples within the buffer, significantly enhancing knowledge retention and utilization efficiency. Our approach provides a flexible solution for dynamic buffer allocation, effectively addressing the varying importance and learning complexity of samples over time, and improves CL performance.[1]

1 Introduction

Continual Learning (CL), the gradual acquisition of new concepts (classes or tasks) without forsaking previous ones, stands as a pivotal capability in machine learning. Various methodologies, including regularization techniques Kirkpatrick et al. (2017); Chaudhry et al. (2018a), architecture-based approaches Mallya & Lazebnik (2018); Hung et al. (2019), and rehearsal-based strategies Chaudhry et al. (2018b; 2019); Aljundi et al. (2019b); Chaudhry et al. (2021); Sun et al. (2022; 2023); Tong et al. (2025), have been explored. However, the central challenge remains in striking a balance between assimilating new concepts (plasticity) and preserving existing knowledge (stability) Kim et al. (2023); Kim & Han (2023).

Although buffer-based methods outperform other approaches by mitigating catastrophic forgetting through the storage and reuse of samples from previously learned concepts in a small buffer, memory constraints make the buffer-filling strategy a critical challenge. While this approach can achieve strong performance Tong et al. (2025), its effectiveness heavily depends on how the limited memory is utilized.

Some buffer-based methods, such as Chaudhry et al. (2019); Lin et al. (2023); Lopez-Paz & Ranzato (2017); Chaudhry et al. (2018b), do not use explicit buffer update policies. For example, Chaudhry et al. (2019); Lin et al. (2023) rely on reservoir sampling, while Lopez-Paz & Ranzato (2017); Chaudhry et al. (2018b) retain the earliest samples from each task. To improve buffer quality, other methods propose sample selection strategies based on gradient informativeness Aljundi et al. (2019b); Jin et al. (2021); Yoon et al. (2021); Tiwari et al. (2022), or on Shapley values and influence functions Shim et al. (2021); Sun et al. (2022; 2023). However, these approaches are typically computationally expensive and slow (see Table 5).

Moreover, prototype-based approaches such as iCaRL Rebuffi et al. (2017), while gradient-free and sometimes effective in improving buffer utility, incur additional computational overhead to compute and update class feature means. Hardest-to-classify-sample strategy Aljundi et al. (2019a) can be beneficial in certain

[1] Code is available in the supplementary materials.

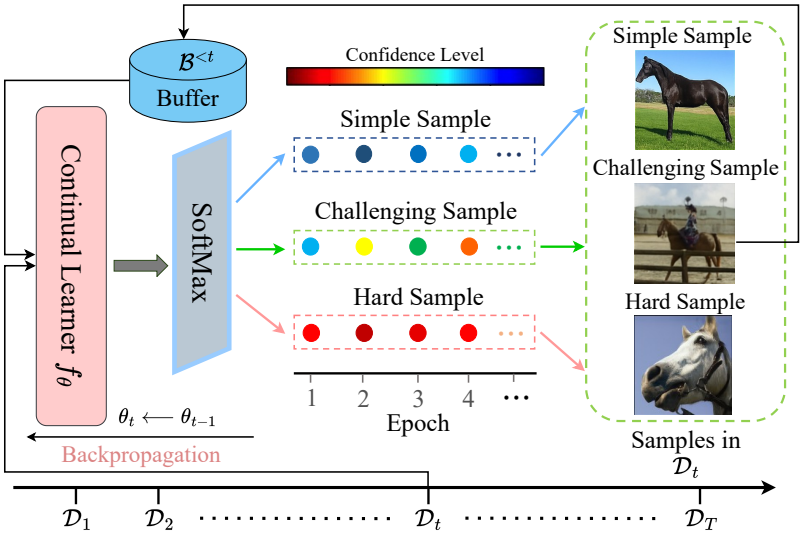

Figure 1: By tracking the target confidence of the main continual learner f_θ during training on each task $\mathcal{D}_t$, our method efficiently prioritizes replaying class samples with the highest confidence fluctuations (i.e., challenging, boundary samples), promoting robust decision boundaries and improving classification margins at minimal cost, unlike prior methods that rely on auxiliary models or gradient computations.

cases, but often risks allocating buffer space to outliers. Importantly, both types of methods tend to overlook samples whose confidence fluctuates over time—a hallmark of low-margin, decision-boundary-critical examples.

A recent approach, CSReL Tong et al. (2025), avoids gradient computations by using forward-pass-only loss differences between a model trained on the coreset and another trained on the full task data. While more efficient than gradient- or influence-based methods, CSReL still requires training auxiliary models per task and performing forward passes per candidate samples, introducing non-trivial computational overhead (see Table 6). This added complexity can limit its practicality in resource-constrained settings. Moreover, because the auxiliary models are trained independently for each task and discarded after selection, they remain decoupled from the main continual learner and do not explicitly capture the informative dynamics of samples over training time—an aspect crucial for identifying boundary-critical or challenging examples.

To address the aforementioned challenges, we propose Confidence-Guided Replay (CGR) (Figure 1), a lightweight and effective method that leverages the confidence dynamics of the main continual learner, enabling a buffer update mechanism at minimal computational cost that does not require gradient computation, influence estimation, or auxiliary training—unlike prior methods—to identify and retain the most informative training samples (i.e., challenging, boundary samples). CGR tracks the variance in the target confidence for each sample across training epochs, targeting those with high variability—indicative of ambiguity and proximity to decision boundaries. CGR is modular and compatible with existing buffer-based methods, improving their performance without altering their core learning processes. Its design aligns with margin-theoretic principles, which emphasize the value of low-margin, uncertain samples for robust generalization Bartlett & Mendelson (2002). By continually reinforcing these high-variance examples, CGR enhances boundary stability and mitigates catastrophic forgetting.

Traditionally, samples with consistently high confidence are viewed as simple (prototypical), while those with persistently low confidence are seen as hard (outliers). In contrast, CGR identifies a critical third category: challenging samples—those with high variance in confidence over time. These are neither confidently learned nor consistently misclassified, but instead indicate decision boundary uncertainty. By prioritizing these samples, CGR selectively reinforces the most informative and challenging instances encountered during training.

The primary contributions of this work are:

1. We address the high computational cost associated with existing rehearsal-based CL methods that rely on complex buffer update policies.

2. We propose a lightweight buffer update strategy based solely on the confidence trajectory of the main model, without requiring auxiliary models, gradients, or additional computations.

3. We demonstrate the effectiveness and efficiency of CGR through extensive experiments on standard CL benchmarks, and show its compatibility as a plug-in to boost the performance of existing CL methods.

2 Related Work

Rehearsal-based CL Methods. Rehearsal-based methods prevent catastrophic forgetting by maintaining a memory buffer of past data and periodically retraining the model with both old and new information. This approach enhances the model's ability to retain the knowledge of previous tasks Chaudhry et al. (2018b; 2019); Aljundi et al. (2019b); Prabhu et al. (2020); Chaudhry et al. (2021); Sun et al. (2022; 2023); Tong et al. (2025).

A-GEM Chaudhry et al. (2018b) leverages previous training data to minimize gradient interference explicitly. They populate the buffer using the earliest samples from each task, ensuring equal task-wise allocation. ER Chaudhry et al. (2019) presents a straightforward rehearsal method that updates memories through reservoir sampling and employs random sampling during memory retrieval. Despite its simplicity, it remains a strong baseline. GSS Aljundi et al. (2019b) approaches updating the memory buffer as a constrained optimization problem, to maximize the diversity of sample gradients within the buffer. GDUMB Prabhu et al. (2020) operates by greedily accumulating new training samples into the buffer, which it subsequently uses to train a model from scratch for testing. HAL Chaudhry et al. (2021) employs the old training samples, which are more susceptible to forgetting, as "anchors" to stabilize their predictions. This method enhances replay by adding an objective that minimizes the forgetting of crucial learned data points. MetaSP Sun et al. (2022) utilizes the Pareto optimum of example influence on stability and plasticity, thereby guiding updates to the model and storage management. SOIF Sun et al. (2023) leverages second-order influences to make more informed decisions about which samples to retain in the buffer. CSReL Tong et al. (2025) introduces a gradient-free coreset selection method using reducible loss (ReL), computed via forward-pass-only loss differences. Its variants include CSReL-CL for per-task selection, CSReL-CL-Prv, which reduces task interference using memory, CSReL-RS for streaming data via reservoir sampling scaled by ReL, and ReL-cmb for knowledge distillation. For fair comparison, we use CSReL-CL-Prv and CSReL-ER (based on CSReL-RS) as baselines.

Data-centric. Data-centric approaches emphasize the quality of data over the complexity of models. Data-centric artificial intelligence involves techniques that aim to enhance datasets, thereby enabling the training of models with fewer data requirements Motamedi et al. (2021); Mazumder et al. (2022). Ignoring the critical importance of data has led to inaccuracies, biases, and fairness issues in real-world applications Mazumder et al. (2022). Previous research reveals that not all samples are equally informative; some contribute substantially more to the learning process than others. This is highlighted in studies such as those by Toneva et al. (2018); Swayamdipta et al. (2020), along with related research by Motamedi et al. (2021); Mazumder et al. (2022). For instance, Toneva et al. (2018) and Swayamdipta et al. (2020) leverage model confidence during training to cleanse the dataset. Specifically, Swayamdipta et al. (2020) categorizes the dataset into easy-to-learn, hard-to-learn, and ambiguous subsets, whereas Toneva et al. (2018) distinguishes between forgettable and unforgettable samples. These methods allow for the assessment of each sample's contribution during training. Building on these approaches, we introduce CGR to identify and retain the most informative (i.e., challenging, boundary) samples in the buffer.

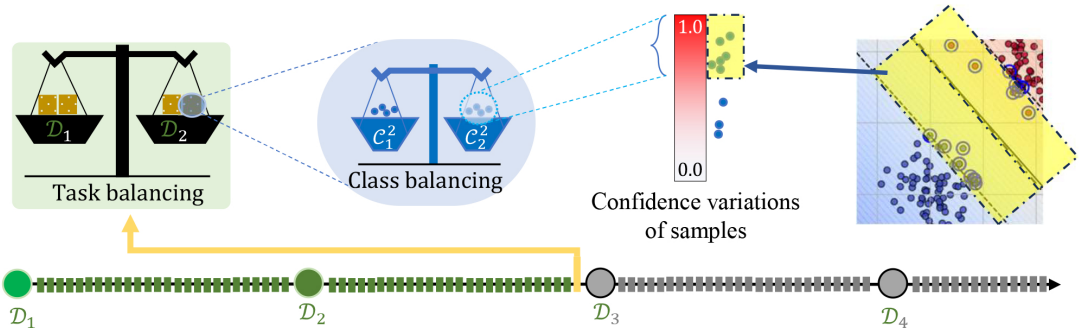

Figure 2: Illustration of the CGR buffer update policy. After training on each task $\mathcal{D}_t$, CGR updates the buffer by selecting the most challenging samples—those exhibiting high temporal variation in target class confidence—while maintaining both class- ($\mathcal{C}$) and task-level ($\mathcal{D}$) balance. The confidence variation heatmap highlights high-variance samples near decision boundaries (yellow region), which are critical for promoting robust generalization. By prioritizing these boundary-sensitive examples, CGR reinforces informative, low-margin instances that are most susceptible to forgetting.

3 Preliminaries

Consider a model f_θ parameterized by θ, which is trained incrementally on a series of tasks $\{1, 2, \ldots, T\}$. At a specific time t, the model processes the task $\mathcal{D}_t = \{(x_i^t, y_i^t), \mathcal{C}^t\}_{i=1}^{N^t}$, where x_i^t represents the input data, y_i^t the corresponding label, N^t the number of samples in $\mathcal{D}_t$, and $\mathcal{C}^t$ indicates the classes specific to task $\mathcal{D}_t$. The buffer $\mathcal{B}$ is used to store a subset of past examples to mitigate catastrophic forgetting while having a limited capacity, denoted as B_{size}. For each task $\mathcal{D}_t$, the model's objective is to learn the new task while retaining knowledge from previous tasks. The model parameters are updated from θ_{t-1} to θ_t by minimizing the loss function over the combined dataset of the current task and the buffer:

$$\theta_t^* = \arg\min_{\theta_t} \mathbb{E}_{(x,y)\sim(\mathcal{D}_t \cup \mathcal{B})} \left[\ell(f_{\theta_t}(x), y) \right], \tag{1}$$

where ℓ represents the loss function.

4 Proposed Method: Confidence-Guided Replay (CGR)

In rehearsal-based CL, buffer updates play a key role in mitigating forgetting. Given the limited size of the buffer $\mathcal{B}$, the model must carefully decide which new samples to store and which old ones to discard. An effective strategy should retain the most valuable examples for future learning.

We propose Confidence-Guided Replay (CGR), a lightweight and modular buffer update policy designed to preserve boundary-critical examples—samples that lie near class decision boundaries and are therefore more susceptible to forgetting. Instead of relying on gradients, auxiliary models, or feature mean (prototype) calculation, CGR leverages a single, low-cost statistic derived from the main model itself: the temporal variance of the model's confidence in its target predictions for each training sample. This confidence trajectory—capturing how predicted target confidence fluctuates across initial training epochs—is used to identify and retain high-variance (i.e., boundary) samples. Specifically, CGR interprets high-variance samples as neither simple (consistently high confidence) nor hard outliers (consistently low confidence), but as challenging and informative decision-boundary samples. By explicitly retaining these ambiguous examples at the end of each task, CGR promotes decision boundary stability, reduces forgetting, and ensures a balanced, representative buffer through class-wise and task-wise selection (see Figures 1 and 2).

Algorithm 1 Confidence-Guided Replay (CGR)

Input: Task sequence $\{\mathcal{D}_1, \ldots, \mathcal{D}_T\}$, class counts $\{|\mathcal{C}^1|, \ldots, |\mathcal{C}^T|\}$, model f_θ, buffer $\mathcal{B}$ with capacity B_{size}, CGR hyperparameter E, epochs number m, current-task and buffer batch size b
Initialize: Empty replay buffer $\mathcal{B} \leftarrow \emptyset$, parameters θ
for $t = 1$ to T **do**
  for epoch $e = 1$ to m **do**
    for $(x_i^t, y_i^t)^b \sim \mathcal{D}_t$ **do**
      $(x_i^{\mathcal{B}}, y_i^{\mathcal{B}})^b \leftarrow \text{RandomRetrieval}(\mathcal{B})$
      $(\mathbf{x}, y) \leftarrow \left\{(x_i^t, y_i^t)^b, (x_i^{\mathcal{B}}, y_i^{\mathcal{B}})^b\right\}$
      if $e \leq E$ **then**
        $\Gamma_e(x_i^t, y_i^t)^b \leftarrow \mathcal{P}\left((y_i^t)^b \mid f_\theta((x_i^t)^b)\right)$
      end if
      $\theta_t^* = \arg\min_\theta \text{CrossEntropy}\left(f_\theta(\mathbf{x}), y\right)$
    end for
  end for
  $\sigma^2(x^t, y^t) = \frac{1}{E}\sum_{e=1}^{E}\left[\Gamma_e(x^t, y^t) - \overline{\Gamma}(x^t, y^t)\right]^2$
  if $t \geq 2$ **then**
    $k_{\text{prune}} \leftarrow \frac{B_{\text{size}}}{(t-1)\,|\mathcal{C}^{t-1}|} - \frac{B_{\text{size}}}{t\,|\mathcal{C}^t|}$
    $\mathcal{B} \leftarrow \mathcal{B} \setminus \left\{(x_{i,c}^{<t}, y_{i,c}^{<t}) \,\middle|\, \sigma^2(x_{i,c}^{<t}, y_{i,c}^{<t}) \in \text{bottom-}k_{\text{prune}}\right\}$
  end if
  $k_{\text{add}} \leftarrow \frac{B_{\text{size}}}{t\,|\mathcal{C}^t|}$
  $\mathcal{B} \leftarrow \mathcal{B} \cup \left\{(x_{i,c}^t, y_{i,c}^t) \,\middle|\, \sigma^2(x_{i,c}^t, y_{i,c}^t) \in \text{top-}k_{\text{add}}\right\}$
end for
Output: Updated parameters θ, and buffer $\mathcal{B}$

Confidence Variance. The target confidence of the model in its predictions is defined as the probability assigned to the target class y_i^t for sample x_i^t during task $\mathcal{D}_t$. Formally, the confidence score is expressed as:

$$\Gamma(x_i^t, y_i^t) = \mathcal{P}\big(y_i^t \mid f_\theta(x_i^t)\big)\,, \tag{2}$$

where f_θ is the model that maps the input x_i^t into class probabilities. During training, misclassified samples fall into two categories: boundary samples and outlier samples. Outlier samples consistently exhibit low confidence, whereas boundary samples display fluctuating confidence levels (Figure 1), reflecting uncertainty in their classification Toneva et al. (2018); Swayamdipta et al. (2020).

This fluctuation can be quantified through the variance of the model's confidence scores computed over the first E epochs of training:

$$\sigma^2(x_i^t, y_i^t) = \frac{1}{E}\sum_{e=1}^{E}\left[\Gamma_e(x_i^t, y_i^t) - \overline{\Gamma}(x_i^t, y_i^t)\right]^2, \tag{3}$$

where the mean confidence is given by

$$\overline{\Gamma}(x_i^t, y_i^t) = \frac{1}{E}\sum_{e=1}^{E}\Gamma_e(x_i^t, y_i^t). \tag{4}$$

Here, $\Gamma_e(x_i^t, y_i^t)$ denotes the model's target confidence for sample (x_i^t, y_i^t) at epoch e, and $\overline{\Gamma}(x_i^t, y_i^t)$ represents the average confidence across E epochs. A higher value of $\sigma^2(x_i^t, y_i^t)$ indicates greater fluctuation in confidence, suggesting that the sample lies near a decision boundary where the model's predictions are less stable.

The full process of CGR is summarized in Algorithm 1, which illustrates how confidence-guided sample selection is integrated into the learning process.

Buffer Structure. To avoid long-tailed memory distributions that distort decision boundaries Cui et al. (2019); Samuel & Chechik (2021); Shi et al. (2023), CGR enforces a balanced allocation across both tasks and classes (Figure 2). This design mitigates over-representation of dominant categories and ensures uniform replay coverage.

Assuming that all tasks $\{\mathcal{D}_1, \ldots, \mathcal{D}_T\}$ contain an equal number of classes and samples, the replay buffer at time t allocates an equal memory share to each task, i.e., $\frac{B_{\text{size}}}{t}$. For a task $\mathcal{D}_t$ with $|\mathcal{C}^t|$ classes, each class receives $\frac{B_{\text{size}}}{t \times |\mathcal{C}^t|}$ capacity. Under this assumption, maintaining class balance inherently ensures task balance. Therefore, CGR does not require explicit task identifiers—only samples and class labels, as in conventional replay-based methods.

During buffer updates, samples with the highest variance within each class are retained, while those with the lowest variance are discarded. Because samples in each class-specific buffer segment are stored in descending order of variance relative to their buffer index, the variance values need not be stored explicitly. Thus, the buffer maintains only the samples and their corresponding labels.

Adaptive Buffer Update and Pruning. To accommodate memory constraints, CGR adopts an *adaptive buffer management strategy* that dynamically prunes less informative samples prior to integrating data from a new task $\mathcal{D}_t$. Specifically, for each previously observed class $c \in \mathcal{C}^{<t}$, the algorithm removes the k_{prune} samples exhibiting the lowest variance from the replay buffer $\mathcal{B}$. The number of samples to be pruned per class is defined as:

$$k_{\text{prune}} = \frac{B_{\text{size}}}{(t-1)\,|\mathcal{C}^{t-1}|} - \frac{B_{\text{size}}}{t\,|\mathcal{C}^t|}. \tag{5}$$

Accordingly, the pruning operation is expressed as:

$$\mathcal{B} \leftarrow \mathcal{B} \setminus \left\{ (x_{i,c}^{<t}, y_{i,c}^{<t}) \,\middle|\, \sigma^2(x_{i,c}^{<t}, y_{i,c}^{<t}) \in \text{bottom-}k_{\text{prune}} \right\}. \tag{6}$$

This selective removal ensures that the buffer retains high-variance samples, which are generally more informative for preserving discriminative decision boundaries and mitigating catastrophic forgetting.

After pruning, the buffer is updated to incorporate samples from the current task $\mathcal{D}_t$. For each class $c \in \mathcal{C}^t$, the k_{add} samples with the highest variance are selected and inserted into the available buffer space:

$$\mathcal{B} \leftarrow \mathcal{B} \cup \left\{ (x_{i,c}^{t}, y_{i,c}^{t}) \,\middle|\, \sigma^2(x_{i,c}^{t}, y_{i,c}^{t}) \in \text{top-}k_{\text{add}} \right\}, \tag{7}$$

where

$$k_{\text{add}} = \frac{B_{\text{size}}}{t\,|\mathcal{C}^t|}. \tag{8}$$

This formulation guarantees *class- and task-balanced memory allocation* while prioritizing the most informative and challenging samples, thereby enhancing the model's capacity for continual adaptation and stable performance across tasks.

4.1 Which Boundary?

As decision boundaries evolve during training, an essential question arises: which boundary should guide sample selection for buffer retention? Should we rely on boundaries influenced by outliers, or prioritize those shaped by challenging, informative samples? Data-centric studies Toneva et al. (2018); Swayamdipta et al. (2020) suggest that early-stage uncertainty is a strong indicator of a sample's long-term contribution to learning.

To investigate this, we analyze CIFAR-10 in a batch training setup—analogous to a single task in Split CIFAR-100. For each sample, we track its target confidence across 4 and 50 epochs and visualize the mean and variance in Figure 3, sub-figures (b.1) and (b.2). These plots reveal how confidence variability highlights boundary-critical samples early in training.

We categorize samples observed in the first $E = 4$ epochs into three groups. *Area 1* (Simple): High confidence, low variance—well-classified and central. *Area 2* (Challenging): Moderate confidence, high variance—near

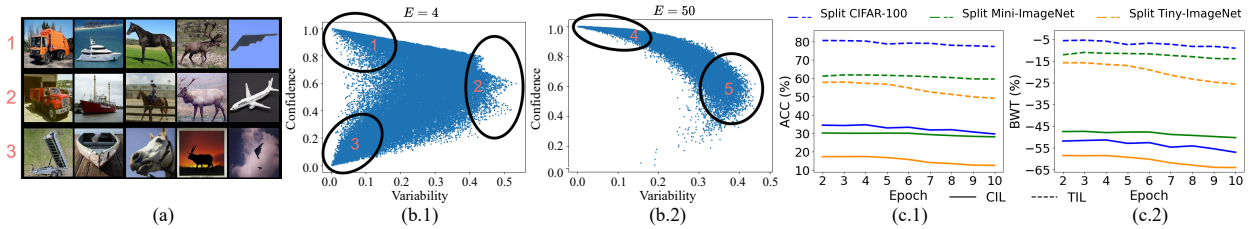

Figure 3: (a) The first, second, and third rows show simple, challenging, and hard samples, respectively, from five classes of CIFAR-10. (b) Visualization of the mean and variance of target confidence for CIFAR-10 samples through epochs 4 and 50. (c) Performance comparison for different values of E in our boundary sample identification, using a buffer size of 1000. CIL and TIL indicate class and task incremental settings, respectively.

boundaries and most informative. *Area 3* (Hard): Low confidence, low variance—rare or atypical, offering little value for boundary refinement.

Sub-figure (a) illustrates representative examples from these areas across five CIFAR-10 classes. Simple samples are easily identifiable, showcasing prominent class attributes or centrally-placed objects (e.g., a horse on grass). Challenging samples exhibit ambiguous traits, such as an airplane against a gray background, deviating from typical class patterns. Hard samples contain rare or atypical features—for instance, a horse showing only its head—and often act as outliers, diverging from the standard sample distribution.

By epoch 50, simple and previously challenging samples consolidate into *Area 4*, while hard samples and remaining challenging cases form *Area 5*. Comparing *Area 2* (early) to *Area 5* (late), we argue that retaining high variance (boundary) samples from early training better supports generalization and decision boundary robustness than late training.

Finally, sub-figures (c.1) and (c.2) assess sensitivity to the epoch window E. We find CGR remains effective for $E = 2$ to 6, beyond which performance declines as hard samples (outliers) transition to boundary (high variance) samples, and challenging to simple ones, reducing their impact on refining the decision boundary.

5 Experiments

5.1 Setup

For our experiments, we extended the Mammoth framework as described in Buzzega et al. (2020); Boschini et al. (2022) to ensure a fair setup for comparing CL methods. All experiments are conducted using a Quadro RTX 8000 equipped with CUDA Version 10.2.

Datasets. Our experiments utilized three benchmarks. The first, Split CIFAR-100, consists of 100 classes split into 10 tasks, each with 10 classes and 32 × 32 pixel images Boschini et al. (2022). The second, Split Mini-ImageNet, divides 100 classes into 5 tasks of 20 classes each, with similarly resized images Sun et al. (2022) (see the ablation study in Section 5.3 for results using the original 84 × 84 resolution). The third, Split Tiny-ImageNet, includes 200 classes split into 10 tasks of 20 classes each, with images resized to 32 × 32 pixels Buzzega et al. (2020). Each dataset has 500 training samples per class, with test samples numbering 100 for Split CIFAR-100 and Mini-ImageNet, and 50 for Split Tiny-ImageNet.

Metrics. We employ two primary metrics to evaluate the performance of CL methods: Average Accuracy (ACC) and Backward Transfer (BWT), as outlined in Lopez-Paz & Ranzato (2017). The ACC metric assesses overall performance by calculating the average accuracy across all tasks after the model has completed training. Conversely, BWT measures the degree of knowledge loss over time by measuring the average forgetting across all previous tasks once training on all tasks has concluded. These metrics are mathematically

Table 1: Results of class-incremental learning: Average Accuracy (ACC, higher is better) and Backward Transfer (BWT, favoring less negative values), averaged over 5 runs. 'Mean' columns within each dataset show averages across buffer sizes, while the 'Mean' column next to the datasets indicates the overall average for each method across all datasets. The best results are bolded, and the second best are underlined. Note: BWT of the GDUMB method was excluded due to computational constraints.

Method	Dataset	Split CIFAR-100				Split Mini-ImageNet				Split Tiny-ImageNet				Mean
	Buffer	500	1000	2000	Mean	500	1000	2000	Mean	500	1000	2000	Mean	
A-GEM	ACC	9.32	9.23	9.25	9.27	14.65	14.69	14.74	14.69	6.30	6.30	6.23	6.28	10.08
	BWT	-85.68	-85.05	-86.07	-85.60	-66.56	-66.79	-66.60	-66.65	-64.18	-64.71	-64.61	-64.50	-72.25
ER	ACC	19.29	24.46	31.94	25.23	18.02	20.78	24.72	21.17	8.90	10.75	13.88	11.18	19.19
	BWT	-75.53	-69.29	-61.18	-68.67	-63.62	-60.07	-55.12	-59.60	-68.53	-67.12	-63.47	-66.37	-64.88
GSS	ACC	22.23	28.84	35.27	28.78	18.48	21.91	28.16	22.85	9.23	11.45	15.52	12.07	21.23
	BWT	-70.26	-61.81	-53.51	-61.86	-61.72	-55.95	-46.80	-54.82	-67.56	-64.05	-58.06	-63.22	-59.97
GDUMB	ACC	10.36	16.35	25.47	17.39	7.09	10.15	14.33	10.52	3.59	5.15	7.53	5.42	11.11
	BWT	–	–	–	–	–	–	–	–	–	–	–	–	–
HAL	ACC	10.23	12.62	16.68	13.18	5.95	7.10	8.59	7.21	2.45	2.58	2.91	2.65	7.68
	BWT	**-55.71**	-53.55	-47.52	-52.26	**-38.88**	**-38.11**	**-36.63**	**-37.87**	**-34.95**	**-34.65**	**-32.17**	**-33.92**	**-41.35**
MetaSP	ACC	21.87	26.23	29.74	25.95	22.92	27.84	33.39	28.05	11.34	14.72	19.22	15.09	23.03
	BWT	-71.01	-65.53	-59.90	-65.48	-60.40	-53.28	-44.64	-52.77	-67.63	-63.03	-56.61	-62.42	-60.23
SOIF	ACC	18.85	26.19	25.36	23.47	18.75	22.85	23.46	21.69	9.10	11.66	11.58	10.78	18.64
	BWT	-71.47	-62.91	-62.11	-65.50	-65.70	-58.65	-54.76	-59.70	-69.63	-65.65	-63.64	-66.31	-63.84
CSReL-ER	ACC	20.29	27.65	36.37	28.10	19.11	23.54	29.22	23.96	9.70	12.10	16.58	12.79	21.62
	BWT	-74.02	-65.41	-55.05	-64.83	-62.96	-56.74	-50.52	-56.74	-68.22	-65.38	-59.71	-64.44	-62.00
CSReL-CL-Prv	ACC	26.61	32.27	35.18	31.35	23.31	28.25	34.54	28.70	11.50	15.57	21.15	16.07	25.38
	BWT	-61.31	-51.37	-44.39	-52.36	-57.60	-48.51	-36.68	-47.60	-65.44	-60.05	-50.83	-58.77	-52.91
CGR	ACC	**28.41**	**34.65**	**41.03**	**34.70**	**24.20**	**30.05**	**36.36**	**30.20**	**12.35**	**17.32**	**23.50**	**17.72**	**27.54**
	BWT	-60.03	**-51.34**	**-41.53**	**-50.97**	-56.39	-47.38	-37.17	-46.98	-64.99	-58.59	-49.42	-57.67	-51.87

defined as $\text{ACC} = \frac{1}{T}\sum_{j=1}^{T}\alpha_{T,j}, \quad \text{BWT} = \frac{1}{T-1}\sum_{j=1}^{T-1}\beta_{T,j}$. Here, $\alpha_{i,j}$ represents the accuracy on the test set held out for task j after the network has been trained on task i. $\beta_{i,j}$ is calculated as $(\alpha_{i,j} - \alpha_{j,j})$, signifying the decrease in model performance on task j due to training on subsequent tasks. T shows the number of tasks. We evaluate these metrics under class incremental learning (CIL) and task incremental learning (TIL) settings, as discussed in previous studies Boschini et al. (2022); Sun et al. (2023).

Baselines. We compared our method against nine rehearsal-based approaches, including A-GEM Chaudhry et al. (2018b), ER Chaudhry et al. (2019), GSS Aljundi et al. (2019b), GDUMB Prabhu et al. (2020), HAL Chaudhry et al. (2021), MetaSP Sun et al. (2022), SOIF Sun et al. (2023), and two variants of CSReL Tong et al. (2025), CSReL-ER, and CSReL-CL-Prv.

Implementation Details. For the implementation and hyperparameters of the baselines, we followed the Mammoth framework Buzzega et al. (2020); Boschini et al. (2022), MetaSP Sun et al. (2022), SOIF Sun et al. (2023), and CSReL Tong et al. (2025), using ResNet18 He et al. (2016) as the backbone model, trained from scratch with cross-entropy loss. We ensured fairness by averaging results over five runs with fixed seeds (0 to 4). Both the current and memory batch sizes were set to 32. The hyperparameter E was empirically set based on buffer size and dataset: for Split CIFAR-100, values of 3, 4, and 5 were used for buffer sizes 500, 1000, and 2000, respectively; for Split Mini-ImageNet, the values were 3, 3, and 4; and for Split Tiny-ImageNet, 3 was used across all buffer sizes. A sensitivity analysis of E is provided in Ablation. We used the SGD optimizer for 50 epochs per task, with learning rates of 0.1 for Split CIFAR-100 and 0.03 for both Split Mini-ImageNet and Split Tiny-ImageNet, applying standard augmentations, including random cropping and random horizontal flipping (see Appendix A.1 for more details).

5.2 Main Results

Class Incremental Learning. Table 1 presents the Average Accuracy (ACC) and Backward Transfer (BWT) results in the class incremental settings. As is evident, our method consistently outperforms other methods in all scenarios for ACC. Regarding BWT, for Split CIFAR-100, we outperform others on average, but for two other datasets, we rank second. However, a closer analysis is warranted. In these cases, the HAL method achieves the best average BWT. Nonetheless, as results show, HAL has the worst average ACC,

Table 2: Results of task-incremental learning: Average Accuracy (ACC, higher is better) and Backward Transfer (BWT, favoring less negative values), averaged over 5 runs. 'Mean' columns within each dataset show averages across buffer sizes, while the 'Mean' column next to the datasets indicates the overall average for each method across all datasets. The best results are bolded, and the second best are underlined. Note: BWT of the GDUMB method was excluded due to computational constraints.

Method	Dataset	Split CIFAR-100				Split Mini-ImageNet				Split Tiny-ImageNet				Mean
	Buffer	500	1000	2000	Mean	500	1000	2000	Mean	500	1000	2000	Mean	
A-GEM	ACC	54.13	54.46	55.86	54.82	32.48	32.56	33.82	32.95	19.24	19.84	19.53	19.54	35.77
	BWT	-35.89	-34.78	-34.29	-34.99	-44.26	-44.46	-42.74	-43.82	-49.80	-49.67	-49.84	-49.77	-42.86
ER	ACC	73.90	78.12	81.69	77.90	50.70	55.25	58.91	54.95	42.15	48.90	54.18	48.41	60.42
	BWT	-15.40	-10.70	-6.97	-11.02	-22.92	-17.42	-13.15	-17.83	-31.73	-25.16	-19.42	-25.44	-18.10
GSS	ACC	72.85	76.75	80.89	76.83	50.23	54.68	59.38	54.76	40.33	47.20	52.80	46.78	59.46
	BWT	-15.91	-12.38	-7.91	-12.07	-23.16	-17.52	-12.36	-17.68	-33.74	-25.99	-20.18	-26.64	-18.80
GDUMB	ACC	32.91	42.73	54.49	43.38	18.15	23.19	29.83	23.72	15.12	18.78	24.51	19.47	28.86
	BWT	–	–	–	–	–	–	–	–	–	–	–	–	–
HAL	ACC	45.83	50.19	55.48	50.50	20.74	24.04	26.65	23.81	15.91	17.60	18.84	17.45	30.59
	BWT	-19.52	-15.96	-10.81	-15.43	-20.69	-17.55	-15.60	-17.95	-20.63	-19.17	-17.09	-18.96	-17.45
MetaSP	ACC	72.55	77.29	80.68	76.84	53.68	58.63	63.19	58.50	44.51	50.89	56.98	50.79	62.05
	BWT	-17.24	-11.79	-7.83	-12.29	-23.32	-17.20	-11.29	-17.27	-31.76	-24.72	-17.96	-24.81	-18.12
SOIF	ACC	70.74	71.82	62.16	68.24	51.02	50.78	46.36	49.39	38.10	37.43	31.17	35.57	51.07
	BWT	-16.36	-16.00	-27.18	-19.85	-26.24	-26.05	-32.03	-28.11	-37.98	-38.56	-45.79	-40.78	-29.58
CSReL-ER	ACC	70.11	75.74	80.03	75.29	53.17	57.50	62.23	57.63	45.55	51.30	56.03	50.96	61.30
	BWT	-19.19	-13.06	-8.63	-13.63	-20.56	-14.87	-10.38	-15.27	-28.57	-22.31	-16.94	-22.61	-17.17
CSReL-CL-Prv	ACC	76.61	78.10	78.34	77.68	**59.66**	**62.03**	**65.28**	**62.32**	**53.76**	**58.31**	**62.44**	**58.17**	66.06
	BWT	-10.15	-8.69	-8.10	-8.98	**-13.25**	**-9.31**	**-5.52**	**-9.36**	**-19.19**	**-14.56**	**-9.76**	**-14.50**	**-10.95**
CGR	ACC	**77.50**	**80.24**	**83.19**	**80.31**	58.18	61.80	64.82	61.60	53.14	57.92	61.91	57.66	**66.52**
	BWT	**-8.83**	**-5.93**	**-2.49**	**-5.75**	-15.39	-11.07	-7.67	-11.38	-20.59	-15.86	-11.35	-15.93	-11.02

indicating that it failed to acquire sufficient knowledge to retain, much less forget, which raises concerns about its overall learning capability. Therefore, a fair comparison should consider both BWT and ACC, as a less negative BWT alone is insufficient if the ACC is low. To this end, we should consider BWT alongside ACC. Notably, although CSReL-CL-Prv incurs higher computational costs (see Table 6), our method consistently outperforms it.

Task Incremental Learning. Table 2 shows results in the task incremental settings. As is evident, for Split CIFAR-100, our method outperforms all others in both ACC and BWT. For the other two datasets, we rank second and perform comparably to CSReL-CL-Prv. On average, across all scenarios, we surpass CSReL-CL-Prv in ACC and achieve a BWT that is nearly identical. However, as shown in Table 6, CSReL-CL-Prv incurs substantial computational costs, limiting its practicality in resource-constrained settings.

Modularity of CGR. To demonstrate the compatibility and modularity of CGR, we integrate it into two representative rehearsal-based methods: PCR Lin et al. (2023), and ER-LAS Huang et al. (2023), both of which employ reservoir sampling. As shown in Table 3, incorporating CGR consistently improves performance across all scenarios in both CIL and TIL settings. The only exception occurs for ER-LAS under the Split CIFAR-100 dataset in the CIL setting with a buffer size of 2000, where the ACC remains nearly identical to the baseline.

Table 3: Performance gains from integrating CGR into PCR and ER-LAS.

Method	Dataset	Split CIFAR-100				Split Mini-ImageNet			
	Setting	CIL		TIL		CIL		TIL	
	Buffer	1000	2000	1000	2000	1000	2000	1000	2000
PCR	ACC	33.03	37.63	68.67	75.90	20.20	24.77	48.82	53.48
	BWT	-48.70	-49.98	-19.59	-13.10	-57.65	-52.41	-22.31	-17.29
PCR-CGR	ACC	**36.83**	**42.24**	**72.48**	**76.94**	**29.22**	**32.25**	**52.87**	**57.49**
	BWT	**-31.46**	**-24.69**	**-15.31**	**-10.57**	**-41.26**	**-36.96**	**-18.29**	**-12.58**
ER-LAS	ACC	37.61	**45.69**	78.84	82.30	28.80	34.73	56.64	60.52
	BWT	-49.48	-37.88	-10.02	-6.58	-46.01	-37.82	-15.09	-11.93
ER-LAS-CGR	ACC	**39.85**	45.62	**80.26**	**82.41**	**34.12**	**39.34**	**61.85**	**65.13**
	BWT	**-43.72**	**-33.43**	**-7.46**	**-4.86**	**-40.30**	**-30.30**	**-11.00**	**-7.25**

Confidence-Based Update Policies. In our method, by epoch E, we measure the variance of target confidence for each sample to identify the most challenging ones within each class, while maintaining balanced class and task distributions in the buffer. Table 4 compares CGR with alternative update strategies: selecting samples with the highest (simplest) and lowest (hardest) mean target confidence across epochs E, as well as uniform sampling—each under class and task balance constraints. As shown, CGR outperforms these alternatives in both CIL and TIL settings.

Table 4: Comparing CGR with three update policies—highest mean confidence (simplest samples), lowest mean confidence (hardest samples), and uniform sampling—under class- and task-balanced conditions in both CIL and TIL settings.

Method	Dataset	Split CIFAR-100				Split Mini-ImageNet				Mean	
	Setting	CIL		TIL		CIL		TIL		CIL	TIL
	Buffer	1000	2000	1000	2000	1000	2000	1000	2000		
Uniform	ACC	32.62	40.79	79.11	82.36	27.53	33.90	58.64	62.24	33.71	70.59
	BWT	-52.74	**-41.53**	-6.74	-3.17	-50.87	-39.51	-15.32	-10.15	-46.16	-8.85
Hardest	ACC	19.72	25.86	65.36	70.33	18.56	22.13	43.44	48.72	21.57	56.96
	BWT	-67.90	-59.02	-22.21	-16.66	-62.30	-55.27	-33.95	-27.01	-61.12	-24.96
Simplest	ACC	32.59	38.99	80.01	82.17	29.35	35.32	60.81	64.17	34.06	71.79
	BWT	-53.30	-44.33	**-5.78**	-3.24	-48.73	-38.74	-12.44	-8.33	-46.28	-7.45
CGR	ACC	**34.65**	**41.03**	**80.24**	**83.19**	**30.05**	**36.36**	**61.80**	**64.82**	**35.52**	**72.51**
	BWT	**-51.34**	**-41.53**	-5.93	**-2.49**	**-47.38**	**-37.17**	**-11.07**	**-7.67**	**-44.36**	**-6.79**

Running Time. In Tables 5 and 6, we report per-hour running times of various methods using a Quadro RTX 8000 GPU. Table 5 compares CGR with other methods on Split CIFAR-100 (buffer size: 2000). As shown, CGR is as fast as ER and outperforms others in runtime efficiency.

Table 5: Running times (hours) for methods on Split CIFAR-100 (buffer size: 2000), using a Quadro RTX 8000 GPU.

Method	A-GEM	ER	GSS	HAL	MetaSP	SOIF	CSReL-ER	CSReL-CL-Prv	CGR
Runtime	1.85h	1.4h	15.56h	2.44h	2.87h	8.13h	2.08h	2.52h	**1.4h**

To further analyze the computational cost of CSReL-CL-Prv, Table 6 extends the comparison to Split Mini-ImageNet and Split Tiny-ImageNet, using ER as the baseline.

As all datasets were resized to 32×32 pixels in our experiments, Split Tiny-ImageNet is approximately twice the size of both Split CIFAR-100 and Split Mini-ImageNet. The results show that CGR and ER maintain consistent runtimes when moving from Split CIFAR-100 to Split Mini-ImageNet. In contrast, CSReL-CL-Prv's cost increases significantly. On Split Tiny-ImageNet, both CGR and ER roughly double in runtime, while CSReL-CL-Prv increases by several times.

CGR matches or surpasses ER in runtime efficiency because it updates the buffer once per task—after completing each task's training—rather than batch-by-batch as in reservoir sampling. As a result, the buffer is empty during the first task, making CGR even faster than ER during early training.

Table 6: Running times (hours) for ER (used as the baseline), CSReL-CL-Prv, and CGR across all three datasets (buffer size: 2000), evaluated on a Quadro RTX 8000 GPU.

Dataset	Method	ER	CSReL-CL-Prv	CGR
Split CIFAR-100	**Runtime**	1.4h	2.52h	**1.4h**
Split Mini-ImageNet		1.4h	5.37h	**1.4h**
Split Tiny-ImageNet		2.8h	11.03h	**2.7h**

5.3 Ablation

Results on Split Mini-ImageNet at the Original Resolution (84 × 84). Table 7 reports results on Split Mini-ImageNet for both CIL and TIL settings, using the original image resolution of 84 × 84 pixels (as opposed to the downsampled 32 × 32) and a buffer size of 1000. This evaluation is designed to assess the impact of image downscaling on model performance. As shown, our method still achieves comparable results to the baselines in terms of both ACC and BWT across CIL and TIL settings. Furthermore, as expected, increasing the resolution to 84 × 84 leads to ACC improvements for all methods.

Table 7: Results on Split Mini-ImageNet at the original 84×84 resolution with a buffer size of 1000. ACC and BWT are reported for both CIL and TIL settings. Higher-resolution inputs improve ACC for all methods, and our approach remains competitive across metrics.

Dataset: Split Mini-ImageNet [84 × 84] — buffer size: 1000											
Setting	**Method**	**A-GEM**	**ER**	**GSS**	**GDUMB**	**HAL**	**MetaSP**	**SOIF**	**CSReL-ER**	**CSReL-CL-Prv**	**CGR**
CIL	***ACC***	16.48	26.57	24.34	13.71	8.65	35.09	29.14	29.19	35.09	35.41
	BWT	-76.84	-65.52	-66.45	-	-45.48	-54.47	-63.29	-63.34	-51.21	-51.28
TIL	***ACC***	40.65	66.00	60.28	29.07	27.75	67.51	65.06	67.95	71.03	70.40
	BWT	-46.63	-16.58	-23.81	-	-22.55	-16.38	-19.48	-15.26	-8.40	-9.44

Final Class and Task Distribution in the Buffer. Table 8 presents an analysis of the final buffer composition for each method, highlighting potential long-tailed memory effects Samuel & Chechik (2021); Shi et al. (2023). At the end of training, we record the number of stored samples per class and per task, and compute the Coefficient of Variation (CV) as $\text{CV} = \left(\frac{\text{Standard Deviation}}{\text{Mean}}\right) \times 100\%$. A CV of 0.00 indicates a perfectly balanced distribution, while higher values reflect increasing imbalance and long-tailed memory bias. Our CGR method achieves a CV of 0.00 for both class and task distributions, demonstrating perfectly balanced memory allocation. Similarly, A-GEM and CSReL-CL-Prv also exhibit task-level CVs of 0.00, indicating ideal task balance. ER, HAL, and CSReL-ER yield CVs close to zero, suggesting near-uniform task distributions. In contrast, GSS, MetaSP, and SOIF display substantially higher CV values, reflecting pronounced task-level imbalance. Regarding class distribution, only our method achieves a CV of 0.00, confirming perfect inter-class balance, whereas all other methods show higher CVs, indicating residual class imbalance within the replay buffer.

Table 8: Coefficient of Variation (CV) for class and task distributions in the buffer after training across different methods, using a buffer size of 1000 with Split CIFAR-100.

Method	**A-GEM**	**ER**	**GSS**	**HAL**	**MetaSP**	**SOIF**	**CSReL-ER**	**CSReL-CL-Prv**	**CGR**
CV of Tasks	0.00	06.80	58.79	06.21	26.96	29.87	04.56	0.00	**0.00**
CV of Classes	29.77	31.43	68.38	34.35	53.85	63.97	37.47	44.25	**0.00**

Hyperparameter Sensitivity. In Figure 3, sub-figures (c.1) and (c.2), we examine the sensitivity of our hyperparameter E, which is integral to our memory update policy. As illustrated, the performance of hyperparameter E remains stable across a range from 2 to 6 on all three datasets, maintaining consistency in both ACC and BWT. Additionally, the results indicate that our approach for identifying boundary samples requires only a few initial training epochs.

6 Conclusion

While rehearsal-based methods have shown strong performance in Continual Learning (CL), their effectiveness is often limited by computational constraints and reliance on expensive buffer update strategies. Our proposed method, Confidence-Guided Replay (CGR), addresses these limitations by introducing a lightweight, forward-pass-only approach that tracks target confidence variability to identify boundary-critical samples. By leveraging the temporal dynamics of sample confidence within the main model, CGR enhances decision

boundary robustness and mitigates forgetting—without incurring additional computational cost. This principled and efficient strategy offers a practical alternative to existing rehearsal policies, paving the way for more scalable and adaptive CL systems.

References

Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11254–11263, 2019a.

Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019b.

Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of machine learning research*, 3(Nov):463–482, 2002.

Matteo Boschini, Lorenzo Bonicelli, Pietro Buzzega, Angelo Porrello, and Simone Calderara. Class-incremental continual learning into the extended der-verse. *IEEE transactions on pattern analysis and machine intelligence*, 45(5):5497–5512, 2022.

Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.

Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 532–547, 2018a.

Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018b.

Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc'Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019.

Arslan Chaudhry, Albert Gordo, Puneet Dokania, Philip Torr, and David Lopez-Paz. Using hindsight to anchor past knowledge in continual learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 6993–7001, 2021.

Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9268–9277, 2019.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Zhehao Huang, Tao Li, Chenhe Yuan, Yingwen Wu, and Xiaolin Huang. Online continual learning via logit adjusted softmax. *arXiv preprint arXiv:2311.06460*, 2023.

Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. *Advances in neural information processing systems*, 32, 2019.

Xisen Jin, Arka Sadhu, Junyi Du, and Xiang Ren. Gradient-based editing of memory examples for online task-free continual learning. *Advances in Neural Information Processing Systems*, 34:29193–29205, 2021.

Dongwan Kim and Bohyung Han. On the stability-plasticity dilemma of class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20196–20204, 2023.

Sanghwan Kim, Lorenzo Noci, Antonio Orvieto, and Thomas Hofmann. Achieving a better stability-plasticity trade-off via auxiliary networks in continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11930–11939, 2023.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.

Huiwei Lin, Baoquan Zhang, Shanshan Feng, Xutao Li, and Yunming Ye. Pcr: Proxy-based contrastive replay for online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24246–24255, 2023.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017.

Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2018.

Mark Mazumder, Colby Banbury, Xiaozhe Yao, Bojan Karlaš, William Gaviria Rojas, Sudnya Diamos, Greg Diamos, Lynn He, Alicia Parrish, Hannah Rose Kirk, et al. Dataperf: Benchmarks for data-centric ai development. *arXiv preprint arXiv:2207.10062*, 2022.

Mohammad Motamedi, Nikolay Sakharnykh, and Tim Kaldewey. A data-centric approach for training deep neural networks with less data. *arXiv preprint arXiv:2110.03613*, 2021.

Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 524–540, 2020.

Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.

Dvir Samuel and Gal Chechik. Distributional robustness loss for long-tail learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9495–9504, 2021.

Jiang-Xin Shi, Tong Wei, Yuke Xiang, and Yu-Feng Li. How re-sampling helps for long-tail learning? *Advances in Neural Information Processing Systems*, 36, 2023.

Dongsub Shim, Zheda Mai, Jihwan Jeong, Scott Sanner, Hyunwoo Kim, and Jongseong Jang. Online class-incremental continual learning with adversarial shapley value. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 9630–9638, 2021.

Qing Sun, Fan Lyu, Fanhua Shang, Wei Feng, and Liang Wan. Exploring example influence in continual learning. *Advances in Neural Information Processing Systems*, 35:27075–27086, 2022.

Zhicheng Sun, Yadong Mu, and Gang Hua. Regularizing second-order influences for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20166–20175, 2023.

Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. Dataset cartography: Mapping and diagnosing datasets with training dynamics. *arXiv preprint arXiv:2009.10795*, 2020.

Rishabh Tiwari, Krishnateja Killamsetty, Rishabh Iyer, and Pradeep Shenoy. Gcr: Gradient coreset based replay buffer selection for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 99–108, 2022.

Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. *arXiv preprint arXiv:1812.05159*, 2018.

Ruilin Tong, Yuhang Liu, Javen Qinfeng Shi, and Dong Gong. Coreset selection via reducible loss in continual learning. In *The Thirteenth International Conference on Learning Representations*, 2025.

Jaehong Yoon, Divyam Madaan, Eunho Yang, and Sung Ju Hwang. Online coreset selection for rehearsal-based continual learning. *arXiv preprint arXiv:2106.01085*, 2021.

A Appendix

A.1 Implementation Details

For the CSReL-ER hyperparameters, on Split CIFAR-100, we set the main continual learner's learning rate to 2e-2, the cross-entropy (CE) loss factor to 1.0, and the selection steps to 100. For the selected subset of the current task, the learning rate and number of training steps were set to 5e-3 and 20, respectively. On Split Mini-ImageNet and Split Tiny-ImageNet, we used a main continual learner learning rate of 3e-2 and kept the CE loss factor at 1.0. The selection steps were set to 250. The learning rate and training steps for the selected subset in these datasets were also set to 5e-3 and 20, respectively.

For the CSReL-CL-Prv hyperparameters, on Split CIFAR-100, we set the main continual learner's learning rate to 2e-2, the loss factor to 4.0, and the selection steps to 40. For the selected subset, the learning rate and number of training steps were set to 5e-3 and 7, respectively. For the holdout model, we set the number of epochs, learning rate, and samples per task from previous tasks to 10, 3e-3, and 200, respectively. On Split Mini-ImageNet and Split Tiny-ImageNet, we used a main continual learner learning rate of 2e-2 and kept the loss factor at 4.0. The selection steps were set to 250. The learning rate and training steps for the selected subset in these datasets were set to 5e-3 and 20, respectively. For the holdout model, we set the number of epochs, learning rate, and samples per task from previous tasks to 20, 3e-3, and 200, respectively.