

Temporal Imbalance of Positive and Negative Supervision in Class-Incremental Learning

Jinge Ma Fengqing Zhu
Purdue University, West Lafayette, Indiana, U.S.A.
{ma859, zhu0}@purdue.edu

Abstract

With the widespread adoption of deep learning in visual tasks, Class-Incremental Learning (CIL) has become an important paradigm for handling dynamically evolving data distributions. However, CIL faces the core challenge of catastrophic forgetting, often manifested as a prediction bias toward new classes. Existing methods mainly attribute this bias to intra-task class imbalance and focus on corrections at the classifier head. In this paper, we highlight an overlooked factor—**temporal imbalance**—as a key cause of this bias. Earlier classes receive stronger negative supervision toward the end of training, leading to asymmetric precision and recall. We establish a temporal supervision model, formally define temporal imbalance, and propose **Temporal-Adjusted Loss (TAL)**, which uses a temporal decay kernel to construct a supervision strength vector and dynamically reweight the negative supervision in cross-entropy loss. Theoretical analysis shows that TAL degenerates to standard cross-entropy under balanced conditions and effectively mitigates prediction bias under imbalance. Extensive experiments demonstrate that TAL significantly reduces forgetting and improves performance on multiple CIL benchmarks, underscoring the importance of temporal modeling for stable long-term learning. Code is available [here](#).

1. Introduction

With the rapid development of deep learning in various visual tasks, models in real-world environments are increasingly exposed to dynamically evolving training data distributions. This paradigm is commonly referred to as Continual Learning (CL) [1, 21, 39]. In particular, in the **Class-Incremental Learning (CIL)** setting, new classes are introduced sequentially as tasks, while the majority of old-class data is no longer accessible in future tasks. Such temporal variation in the training distribution brings the core challenge of CIL: catastrophic forgetting [10, 26]. Specifically, when new tasks introduce a large number of new-class sam-

ples, the model’s accuracy on earlier classes drops dramatically. Many studies have indicated that one important cause of this phenomenon is the model’s prediction bias toward new classes [6, 49].

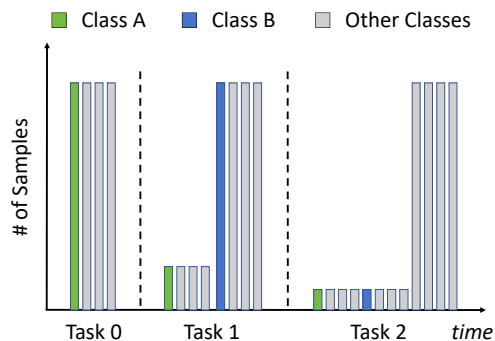


Figure 1. Illustration of temporal imbalance, which differs from class imbalance. In Task 2, although old classes A and B have the same number of training samples, the earlier old class suffers from more severe forgetting than the later one.

To alleviate this issue, many methods treat it as a consequence of class imbalance between new and old classes within each task. As a result, numerous CIL approaches borrow ideas from long-tailed learning. In general, existing methods for reducing prediction bias toward new classes can be roughly grouped into three categories. **(1) Balanced fine-tuning**, which retrains the classifier on a small balanced validation set to remove the bias in classifier heads [49]. **(2) Prototype-based classifier heads**, which replace the standard fully-connected(FC) head with class prototypes in the feature space and classify samples by their similarity (such as cosine) to these prototypes, thereby reducing the dominance of new-class weights [29]. **(3) Output layer calibration**, which directly adjusts the logits or classifier’s parameters through weight alignment or normalization [50]. A more detailed literature review is provided in Sec. 2.

Existing techniques provide effective remedies, however, they mainly operate at the classifier head and explain bias

only from the perspective of class imbalance. In this paper, we argue that attributing the prediction bias toward new classes solely to class imbalance between old and new classes is an oversimplification. For example, as shown in Fig 1(a), in Task 2, suppose that classes A and B appear with the same frequency, so there is no class imbalance between them. Yet, if most positive samples of class A were provided in Task 0 while those of class B were mainly concentrated in Task 1, then class A will have experienced stronger negative supervision at the beginning of Task 2. This demonstrates that even when classes appear balanced in the current task, differences in the historical distribution of positives still lead to unequal supervision biases between old classes. As a consequence, even though there is no class imbalance among these old classes in the last task, earlier old classes generally suffer stronger negative supervision pressure, often resulting in higher precision but lower recall (see Sec. 4.2 for a detailed analysis). In addition, the temporal order of training data in CIL introduces a systematic bias, which affects all model parameters, not limited to the classification head.

To address these issues, and to distinguish temporal imbalance from intra-task class imbalance, we model CIL as a problem of temporal imbalance between positive and negative supervision across the entire training process. We propose a time-order-sensitive loss function, Temporal Adjusted Loss (TAL). Specifically, TAL applies a time-decayed memory kernel to convolve the positive and negative supervision sequences of each class, thereby maintaining a vector Q that characterizes the temporal positive supervision strength. Based on Q , TAL adaptively reweights the negative supervision in the cross-entropy loss, ensuring that the sensitivity to negative supervision is proportional to the class’s temporal positive supervision strength. In this way, TAL dynamically allocates negative pressure according to each class’s current temporal status.

In summary, our main contributions are as follows:

- We establish a temporal supervision model for CIL, define the problem of temporal imbalance in CIL, and propose Temporal Adjusted Loss (TAL) to mitigate it.
- We provide both theoretical and empirical analyses of TAL, derive the steady-state properties of the vector Q , and further introduce a frequency-aligned scaling factor to ensure optimization stability.
- We conduct extensive experiments, demonstrating that TAL effectively reduces forgetting across multiple scenarios.

2. Related Works

2.1. Class-Incremental Learning

Class-Incremental Learning (CIL) is one of the core scenarios of continual learning. Existing methods mainly focus on

alleviating catastrophic forgetting: replay-based [7, 30, 38] approaches maintain an exemplar memory to revisit old classes, regularization-based approaches [21, 23, 24] constrain parameter updates according to their importance to reduce the damage to old knowledge, and dynamic architecture approaches [50, 52, 53] expand the network capacity to adapt to new tasks. As discussed in the introduction, to mitigate the prediction bias towards new classes, three common techniques have been proposed as specific strategies within these approaches. **Balanced fine-tuning** re-trains the classifier on a small balanced validation dataset to reduce bias [9, 43, 49], but it introduces additional training time overhead. Second, **Prototype-based classifiers** compute classification scores based on the similarity between a sample and each class prototype, rather than relying on raw logits, thereby alleviating the dominance bias of new-class weights [8, 18, 29, 55]. However, these methods operate only at the classifier head and fail to address the systematic bias caused by representation drift in the backbone. Third, **output layer calibration** directly rescales or aligns the logits or classifier weights [3, 44, 50–53]. Although lightweight, it is a post-hoc fix only in the fully-connected layers and could not be applied to cosine-based classifiers.

Overall, these methods generally attribute prediction bias to the class imbalance between new and old classes and only correct the classifier head. As a result, they cannot alleviate the heterogeneous forgetting degrees across different old classes, nor do they address biases beyond the classifier head. In contrast, our loss function performs temporal modeling of the supervision signals for each class, enabling fine-grained differentiation across classes and systematically correcting the biases that CIL introduces into the overall model.

2.2. Temporal Modeling in Data Streams

When dealing with data streams, many fields rely on temporal modeling mechanisms to explain the phenomenon that “recent data has a greater impact on the present state.” In time series forecasting, exponential smoothing captures trends by assigning decayed weights to historical observations and has been widely applied in economics and financial forecasting [17, 19, 20, 47]. For example, the Holt–Winters seasonal exponential smoothing method handles time series with both trend and seasonality. In reinforcement learning, eligibility traces and TD(λ) [32, 36, 37], as well as their modern extensions such as Retrace(λ) [28] and Expected Eligibility Traces [40], use **exponential decay** to assign higher “weight” to recent state–action pairs, thereby addressing the problem of delayed rewards. In online learning and concept drift detection, sliding windows and exponential decay are also widely adopted [2, 11, 12, 14]. For instance, ADWIN [4] detects distributional changes in data streams by dynamically adjusting the window size, au-

tomatically expanding or shrinking it in response to drift. In recommender systems, time-aware collaborative filtering methods incorporate temporal decay or windowing on users’ historical behaviors to better model evolving user interests, such as the use of piecewise decay functions for rating prediction [41, 48].

In summary, these approaches share the intuition that recent data should receive larger weights than distant data. However, they are not primarily designed for class-incremental learning. To the best of our knowledge, there has been limited explicit temporal modeling of positive–negative supervision at the loss level in CIL; we take a step in this direction by formalizing temporal imbalance and proposing a remedy for prediction bias.

3. Preliminaries

3.1. CIL Settings

In CIL, we can denote the training dataset for task t as

$$\mathcal{D}_t = \bigcup_{k \in K_t} \{(x_{t,k,i}, y_{t,i})\}_{i=1}^{n_{t,k}} \quad (1)$$

where $x_{t,k,i}$ and $y_{t,i}$ represent the i -th image and its corresponding label for class k in task t , respectively. The value $n_{t,k}$ denotes the number of samples for class k in task t , and K_t is the set of all class indices present in the training data of task t (including replay exemplars if any).

Assume the model parameters are θ and the model has learned the first T tasks and is now learning task $T + 1$. Since the datasets $\bigcup_{t=1}^T \mathcal{D}_t$ for all previous tasks are no longer available, the model can only adapt its existing parameters to fit \mathcal{D}_{T+1} , which typically leads to **catastrophic forgetting**—the model’s performance on new tasks improves while its accuracy on earlier tasks deteriorates significantly.

3.2. Positive and Negative Supervision

For a single sample (x, y) , let $z \in \mathbb{R}^C$ denote the pre-softmax output logits of the classifier, with z_k the k -th logit and C the total number of classes. The cross-entropy (CE) loss can be written as

$$\begin{aligned} \ell_{\text{CE}}(y, z) &= -\log \frac{e^{z_y}}{e^{z_y} + \sum_{k \neq y} e^{z_k}} \\ &= \underbrace{\log \left(e^{z_y} + \sum_{k \neq y} e^{z_k} \right)}_{\text{negative supervision}} - \underbrace{z_y}_{\text{positive supervision}} \end{aligned} \quad (2)$$

Thus, to minimize the ℓ_{CE} for x , as shown in Eq. 2, the “positive supervision” pushes up the true-class logit, while the “negative supervision” corresponds to the softmax denominator in which all classes exert suppressive pressure.

In this sense, for a given class, its positive samples are those belonging to the class (providing positive supervision), while its negative samples are those not belonging to it (providing negative supervision).

4. Methodology

This section introduces how we interpret and rectify the model prediction bias in CIL from the perspectives of temporal order and supervision polarity. We introduce a temporal modeling mechanism to track, for each class k , its **temporal positive supervision strength** Q_k , which characterizes the extent to which the class has been reinforced by positive samples in recent steps. Q_k is recursively updated through an exponentially decaying memory kernel function $f[n]$, reflecting the fact that more recent samples have a stronger influence on the current supervision polarity of that class. Based on Q_k , we propose the **Temporal-Adjusted Loss (TAL)**, which dynamically adjusts the sensitivity to negative supervision according to each class’s current Q_k : for old classes that have recently lacked positive supervision (with smaller Q_k), their negative supervision weights are automatically attenuated; whereas for new classes that have been sufficiently reinforced (with larger Q_k), the sensitivity to negative supervision is preserved. In addition, we introduce a **frequency alignment coefficient** α , which ensures that, for classes with balanced temporal and categorical distributions, TAL degenerates into the standard cross-entropy loss.

4.1. Temporal Modeling of Supervision

Temporal Positive Supervision Strength Q . We consider a simple case where one sample arrives per step (i.e., unit batch size). For a given class k , at step N , we denote its supervision polarity sequence $\{a_k[n]\}_{n=0}^{N-1}$ by:

$$a_k[n] = \begin{cases} +1, & \text{if the sample at step } n \text{ is positive to class } k, \\ -1, & \text{if the sample at step } n \text{ is negative to class } k. \end{cases} \quad (3)$$

To model temporal effects, we introduce a general decay memory kernel function $f[\cdot]$, where $f[n]$ ($n \in \mathbb{N}_0$) represents the residual influence on the current step from a sample which was seen n steps ago (equivalently, the influence of a sample after n steps).

Then, at step N , the temporal positive supervision strength vector $Q[N]$ is defined as

$$\begin{aligned} Q[N] &:= (Q_1[N], Q_2[N], \dots, Q_C[N]), \\ Q_k[N] &= \sum_{n=0}^{N-1} f[N-1-n] a_k[n] = (f * a_k)[N-1]. \end{aligned} \quad (4)$$

where $(f * a_k)$ denotes the discrete convolution between $f[\cdot]$ and the supervision sequence $a_k[\cdot]$.

Intuitively, our defined $Q_k[N]$ measures the effective positive supervision of class k at step N : positive samples increase Q_k , while negative samples decrease it. Naturally, $f[n]$ is defined to satisfy $f[n] \geq 0$ and $f[n+1] \leq f[n]$ for any n , reflecting the fact that more recent training samples have stronger impacts on $Q_k[N]$ due to model’s forgetting. The specific choice of $f[\cdot]$ will be discussed in Sec. 4.3. A larger Q_k indicates stronger positive supervision to class k at the current step, whereas a smaller Q_k implies that positive supervision for the class is relatively lacking.

4.2. Temporal Imbalance

Consider two classes A and B observed over N steps with the *same* total number of positive samples. Up to step n , denote the **number of cumulative positive samples** of class k as $S_k[n]$.

Theorem 1 (Temporal imbalance under equal sample counts). *Consider two classes A and B with the same total number of positive samples. Since $f[n]$ is monotonically decreasing, it follows that*

$$\begin{aligned} S_A[n] &\geq S_B[n] \quad \forall n = 0, 1, \dots, N-1 \\ \implies Q_A[N] &\leq Q_B[N]. \end{aligned} \quad (5)$$

with strict inequality if $\exists n$ such that $S_A[n] > S_B[n]$ and $f[n]$ is strictly decreasing.

The proof of Thm. 1 is deferred to the Appendix.

Interpretation. Thm. 1 formalizes temporal imbalance that is independent of class imbalance: with the same number of positive samples M , the class whose positives occur later in time will attain a larger $Q_k[N]$ at the end of training. In the CIL setting, an early class A attains a smaller $Q_A[N]$ than a late class B , as shown in Fig. 2(a), i.e., $Q_A[N] < Q_B[N]$. Intuitively, a smaller $Q_k[N]$ means that negative supervision dominates at step N , which suppresses the logits of class k and makes predictions for class A concentrate on high-confidence samples, yielding higher precision but lower recall. Conversely, a larger $Q_B[N]$ indicates stronger positive reinforcement, admitting more less confident cases, which improves recall at the cost of precision. As shown in Fig. 2(b), the intuitive effect of temporal imbalance is the asymmetry between precision and recall: at the end of training, earlier classes (e.g., class A) have smaller Q values, leading to higher precision but lower recall, whereas later classes (e.g., class B) have larger Q values, resulting in lower precision but higher recall. As illustrated in Fig. 2(c), our empirical studies demonstrate that this precision–recall asymmetry induced by temporal imbalance is a pervasive phenomenon across various existing CIL methods.

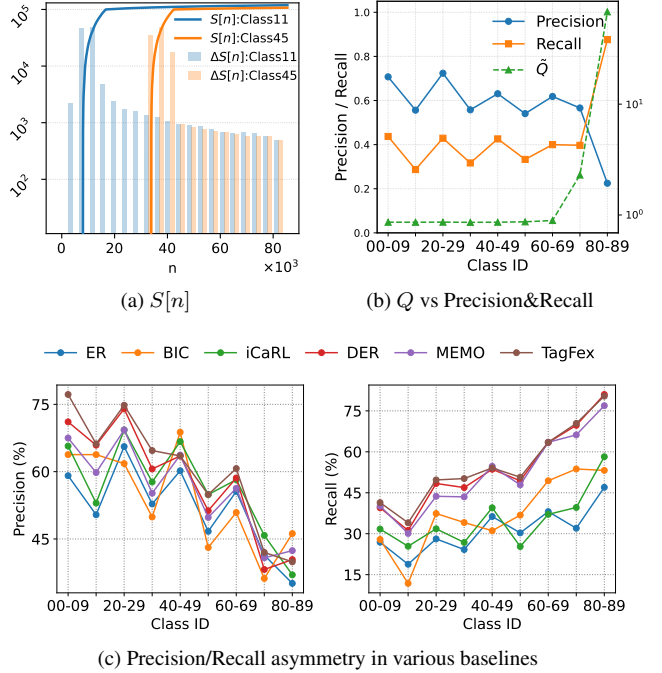


Figure 2. (a) **An example of Temporal Imbalance:** it illustrates, in the CIL setting, how the cumulative number of positive samples $S[n]$ for two classes evolves as the number of training batches n increases. Since the positive samples of class 11 appear earlier than those of class 45, the result shows that $S_{11}[n] \geq S_{45}[n]$. (b)&(c) **Effects of Temporal Imbalance:** (b) Shows, at the end of training, the correlation between the Precision/Recall of each class and its Q value when ER is trained with the cross-entropy loss. For clearer visualization, the normalized value $\tilde{Q} = Q - Q_{\min}$ is used, where Q_{\min} denotes the minimum Q across all classes. (c) Demonstrates that, at the end of training, the Precision–Recall asymmetry induced by temporal imbalance is a common phenomenon across different CIL methods. Note that all classes shown in the figure are old classes, so there is no class imbalance in the final task.

4.3. Choice of the Memory Kernel

Inspired by the exponential temporal decay mechanism in prior work [32], we adopt an exponential decay memory kernel in this paper.

$$f[n] = \lambda^{n+1}, \quad 0 < \lambda < 1. \quad (6)$$

where we refer to λ as the **memory parameter**, which controls how fast the contribution of positive or negative supervision from past samples to $Q_k[N]$, defined in Eq. 4, decays over time. A larger λ means we suppose the $Q_k[N]$ retains stronger memory of earlier supervision’s polarity, whereas a smaller λ implies weaker memory.

Once the hyperparameter λ is fixed, we can derive an asymptotic upper bound Q_{\max} for $Q_k[N]$. Consider the case when $N \rightarrow \infty$ and all training samples belong to class

k (i.e., $a_k[n] \equiv 1$). By the definition of $Q_k[N]$ in Eq. 4, we always have $Q_k[N] < Q_{\max}$ and

$$\begin{aligned} Q_{\max} &= \lim_{N \rightarrow \infty} Q_k[N] = \lim_{N \rightarrow \infty} \sum_{n=0}^{N-1} \lambda^{N-n} a_k[n] \\ &= \sum_{m=1}^{\infty} \lambda^m = \frac{\lambda}{1-\lambda}. \end{aligned} \quad (7)$$

Once the exponential decay memory kernel is fixed, the computation of Q can be viewed as a Markov chain: $Q_k(N+1)$ depends only on $Q_k[N]$ and the current supervision $a_k[N]$, without explicitly storing the entire sequence $\{a_k[n] \mid n = 0, \dots, N-1\}$. Specifically,

$$\begin{aligned} Q_k(N+1) &= \sum_{n=0}^N \lambda^{N+1-n} a_k[n] \\ &= \lambda \left(\sum_{n=0}^{N-1} \lambda^{N-n} a_k[n] \right) + \lambda a_k[N] \\ &= \lambda (Q_k[N] + a_k[N]). \end{aligned} \quad (8)$$

This recursive form reduces the calculation of $Q_k[N]$ from $\mathcal{O}[N]$ to $\mathcal{O}(1)$ time and memory complexity, since only $Q_k[N]$ and $a_k[N]$ are required.

4.4. Our Method: Temporal-Adjusted Loss

From the above analysis, we observe that temporal imbalance causes old classes to receive excessive negative supervision at the end of training, leading to model’s bias toward new classes. To address this issue, we design **Temporal-Adjusted Loss (TAL)**. Our motivation is to dynamically adjust the sensitivity to negative supervision from the perspective of the loss function, based on each class’s current $Q_k[N]$ value. In this way, classes with insufficient positive supervision are protected from excessive negative pressure, while classes with sufficient positive supervision remain highly sensitive to negative supervision.

TAL Function. For a single training sample (x, y) with logits $z \in \mathbb{R}^C$, given the current temporal positive supervision strength $Q[N]$, we define its TAL loss as:

$$\begin{aligned} \ell_{\text{TAL}}(y, z, Q[N]) &= -\log \left(\frac{e^{z_y}}{e^{z_y} + \alpha \sum_{k \neq y} w(Q_k[N]) e^{z_k}} \right), \\ \text{where } w(Q_k[N]) &= \left(\frac{Q_k[N]}{Q_{\max}} \right)^r \end{aligned} \quad (9)$$

By comparing Eq. 2 with Eq. 9, we observe that besides the true label y and the model prediction z , TAL additionally incorporates the current $Q[N]$, thereby introducing temporal information. Specifically, the true class e^{z_y} directly participates in the loss calculation, which is identical

to CE loss. For non-true classes, the e^{z_k} ($k \neq y$) is multiplied by a weight function $w(\cdot)$ positively correlated with $Q_k[N]$, and then further scaled by a frequency alignment parameter α before being included in the loss. The exponent r ($r > 0$) in $w(\cdot)$ controls the steepness of this weighting: larger r makes $w(Q_k[N])$ grow more sharply with $Q_k[N]$. For a class k ($k \neq y$), the closer $Q_k[N]$ is to the upper bound Q_{\max} , the larger $w(Q_k[N])$ becomes, making it more sensitive to negative supervision. Conversely, the closer $Q_k[N]$ is to 0, the less sensitive it is to negative supervision. The range of $Q_k[N]$, the range of $w(\cdot)$, and the derivation of α will be discussed later.

Update of $Q[N]$. With the introduction of the weight function $w(\cdot)$, the recursive update of Q_k in Eq. 8 is modified: the effect of negative supervision is attenuated, while the strength of positive supervision remains unchanged. The update rule is

$$Q_k[N+1] = \lambda \left(Q_k[N] + \begin{cases} a_k[N], & a_k[N] = +1, \\ w(Q_k[N]) a_k[N], & a_k[N] = -1, \end{cases} \right). \quad (10)$$

This ensures that the update of Q is consistent with the attenuation of negative supervision in TAL.

4.5. Parameter Analysis

Boundary conditions of Q and $w(\cdot)$. From Eq. 7 and 10, the temporal positive supervision strength satisfies

$$0 \leq Q_k[N] < Q_{\max} = \frac{\lambda}{1-\lambda}. \quad (11)$$

The lower bound is attainable (e.g., when class k receives no positive supervision), while the upper bound is not attained for $0 < \lambda < 1$. Since $Q_k[N] \in [0, Q_{\max})$, according to the definition of $w(Q_k[N])$ in Eq. 9, we have $0 \leq w(Q_k[N]) < 1$. At the lower bound $Q_k[N] = 0$, we have $w = 0$, meaning class k is completely insensitive to negative supervision because of the lack of recent positive samples, effectively screened out in the loss computation. At the upper limit $Q_k[N] \rightarrow Q_{\max}$, we have $w \rightarrow 1$, recovering full sensitivity to negative supervision. Formal proofs based on the recursion in Eq. 10 are deferred to Appendix.

Frequency alignment parameter α . The role of α is to ensure that, on a temporally uniform (e.g., randomly shuffled) and class-balanced dataset with C classes (each class has prior $p = 1/C$), TAL degenerates to the vanilla cross-entropy loss. In other words, in the absence of temporal imbalance, the negative supervision weight coefficient for each class k on their negative samples should revert to 1. Specifically, we require $\alpha \cdot w(Q^*) = 1$, where Q^* denotes the steady-state temporal positive supervision strength under balanced conditions.

Taking expectation in Eq. 10 and imposing the steady-state condition $Q_k[N+1] = Q_k[N] = Q^*$ yields the calibration equation

$$\left(1 - \frac{1}{C}\right)(x^*)^r + x^* - \frac{1}{C} = 0, \quad \alpha = \frac{1}{(x^*)^r}, \quad (12)$$

where $x^* = Q^*/Q_{\max}$. This equation admits a unique solution $x^* \in (0, 1)$ (**proof in Appendix**). Notably, α is solely determined by the number of classes C and the exponent r .

Summary of parameters. In summary, the temporal modeling introduces several parameters: λ is the memory parameter in Eq. 6, modeling the decay speed of past supervision; $Q_k[N]$ in Eq. 4 tracks the temporal positive supervision strength for class k ; $w(Q_k[N])$ in Eq. 9 rescales the negative supervision to class k based on $Q_k[N]$, with its steepness controlled by exponent r ; and α in Eq. 12 is the frequency alignment parameter that calibrates the overall balance between positive and negative supervision under temporally uniform and class-balanced conditions. Among these, **only λ and r are hyperparameters** to be chosen by the practitioner, while $Q_k[N]$, $w(Q_k[N])$, and α are uniquely determined once λ and r are specified.

5. Experiments

In this section, we conduct two groups of experiments. In Sec. 5.2, we integrate TAL into both classical and most recent CIL methods on three benchmark datasets to observe the improvements brought by TAL. In Sec. 5.3, we perform ablation studies on CIFAR-100 to investigate the effects of hyperparameters λ , r . For all methods, the rehearsal samples are selected using the herding strategy [29].

5.1. Experimental setup

Datasets. We conduct experiments on CIFAR-100 [22], ImageNet-100 [31], and Food101 [5]. CIFAR-100 consists of 50,000 training images and 10,000 test images across 100 classes, with each image being a small 32×32 color image. ImageNet-100 is a 100-class subset of ImageNet, containing 1,300 training samples and 50 test samples per class. Food101 includes 101 food categories with 75,750 training and 25,250 test images.

Implementation Details. We conduct experiments using a public class incremental learning toolboxes PyCIL [54]. To ensure reproducibility, We provide performance results with standard deviations computed over 5 runs. For model backbones, we use ResNet32 on CIFAR-100, and ResNet18 on ImageNet-100 and Food101. All other hyperparameters for the CIL methods and the optimizer are left at the official codebase’s default settings.

5.2. Results and Visualizations

We adopt iCaRL [29], FOSTER [44], DER [50], MEMO [53], and TagFex [52] as baseline methods. In Sec. 5.2, the replay buffer size is **fixed at 20 samples per class across all methods**. Additional results under a fixed total memory budget of 2000 samples are provided in the appendix for all baselines and their TAL variants. As shown in Tab. 1 and Fig. 4, TAL consistently improves the performance of various baselines across different task configurations, regardless of their complexity. Notably, after applying TAL, even the most basic baseline, iCaRL, surpasses more advanced methods such as FOSTER and MEMO in several setups, including CIFAR-100 (20-task), ImageNet-100 (10-task), and Food101 (10-task).

To better illustrate the reasons behind the performance improvements brought by TAL, we visualize the results of four baseline methods (iCaRL, DER, MEMO, and TagFex) on ImageNet-100 under a fixed learning order. Fig. 3 shows the class accuracy curves over class IDs, arranged according to the learning sequence, after the completion of the final task. It can be observed that even with weight alignment, the recall (i.e., class accuracy) of earlier classes remains significantly lower due to the lack of temporal modeling, while relatively newer classes achieve much higher recall. After introducing TAL, however, the model no longer applies the same level of protection to earlier and more recent old classes with equal numbers of exemplars. Instead, it adjusts its prediction bias more reasonably according to the temporal order of classes.

In Fig. 5, we illustrate the impact of TAL on the feature space of the backbone network. We use UMAP [27] to visualize the spatial distribution of feature vectors from five classes sampled across different tasks of iCaRL. It can be observed that earlier classes (such as Class 3 and Class 11) are more prone to being mixed with other classes in the feature space, and their regions are gradually occupied by newer classes (such as Class 38). After introducing TAL, this phenomenon is mitigated. This indicates that TAL not only provides corrections at the classifier level but also has effects that are not limited to linear classification heads.

5.3. Ablation Studies

We investigate the effect of TAL’s two hyperparameters—the memory parameter λ and the steepness exponent r —on model performance, using **Experience Replay (ER)** as the baseline. All experiments are conducted on the CIFAR100&10-task benchmark with a ResNet32 backbone under different replay buffer sizes.

Theoretically, the memory parameter λ represents the estimated decay rate of supervision over time, which is closely related to both model architecture and dataset complexity. An excessively large or small λ can lead to biased loss estimation. Tab. 2 reports the grid search results

Table 1. Average accuracy (A_{Mean}) and last accuracy (A_{Last}) under 10-task and 20-task settings across three datasets. “+ Ours” indicates the baseline method with the Temporal-Adjusted Loss applied. **All methods use a replay buffer of 20 exemplars per class.**

Method	CIFAR-100				ImageNet-100				Food101	
	10-task		20-task		10-task		20-task		10-task	
	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}
iCaRL [29]	58.76±0.12	45.39±0.22	55.98±0.14	43.20±0.24	43.71±0.12	24.38±0.24	33.45±0.14	17.62±0.28	63.46±0.12	50.13±0.22
+ Ours	60.82±0.10	47.36±0.18	58.68±0.12	44.79±0.20	52.19±0.10	32.78±0.20	42.54±0.12	23.90±0.24	68.57±0.10	56.03±0.18
FOSTER [44]	62.04±0.42	49.30±0.64	56.99±0.46	47.34±0.66	48.39±0.46	20.88±0.70	35.67±0.48	15.28±0.72	63.88±0.46	52.60±0.66
+ Ours	64.80±0.40	52.76±0.60	58.21±0.42	49.28±0.64	51.49±0.42	23.38±0.66	38.07±0.46	19.72±0.70	64.61±0.42	53.90±0.60
MEMO [53]	61.68±0.16	51.70±0.30	56.98±0.18	48.93±0.34	47.01±0.16	34.98±0.34	34.73±0.18	22.78±0.36	62.25±0.16	54.75±0.30
+ Ours	63.81±0.12	52.85±0.28	58.04±0.16	49.86±0.30	51.05±0.12	37.28±0.30	38.76±0.16	26.16±0.34	64.73±0.12	55.83±0.28
DER [50]	63.53±0.19	50.75±0.35	57.47±0.17	50.34±0.31	52.25±0.25	40.28±0.49	43.95±0.19	25.82±0.35	69.20±0.23	61.42±0.47
+ Ours	66.33±0.17	53.82±0.31	60.34±0.13	53.49±0.29	54.57±0.23	42.62±0.43	47.80±0.17	29.79±0.31	71.58±0.19	63.98±0.41
TagFex [52]	65.97±0.20	55.99±0.41	59.55±0.20	51.48±0.41	54.73±0.35	41.70±0.59	47.32±0.23	31.36±0.35	71.55±0.31	66.67±0.53
+ Ours	68.68±0.18	57.91±0.37	62.34±0.18	54.70±0.37	57.05±0.31	43.01±0.55	51.08±0.19	35.78±0.31	74.57±0.29	68.64±0.49




Figure 3. This figure shows the performance of different baselines on each class when Task id = 3 under the 10-task setting on ImageNet-100. It can be observed that earlier classes exhibit lower average class accuracy (i.e., recall), similar to Fig 2(c). TAL reduces the sensitivity of old classes to negative supervision while increasing that of new classes. As a result, TAL improves the recall of most old classes at the cost of lowering the recall of new classes slightly.

over λ and r , with a fixed replay buffer size of 2000 and an equal number of replay samples allocated to each old class. We observe that the optimal λ in this experimental setting is 0.995. Moreover, TAL consistently outperforms the cross-entropy loss across a broad hyperparameter range, demonstrating its stability and robustness to hyperparameter choice. Fig. 6 shows the forgetting curves under replay buffer sizes of 1000 and 500, with $\lambda = 0.999$ fixed for all subplots. TAL consistently raises the forgetting curves, in-




Figure 4. Forgetting curves with representative CIL baselines and datasets.




Figure 5. This figure illustrates the impact of applying TAL on the feature space of iCaRL. We visualize the distribution of test images from several new and old classes in the ResNet18 feature space on ImageNet100 at Task ID = 3.

dicating more stable retention over time.

The value of r controls the relationship between the sensitivity to negative supervision and temporal positive supervision across classes. A larger r makes the weight function $w(Q_k[N])$ sharper as $Q_k[N]$ increases, leading to stronger negative supervision on new classes, which enhances the protection of old classes but also limits the recall of new classes. Fig. 7 illustrates the impact of different r values on both new and old classes at the end of the last task. The

Table 2. Ablations on λ and r . We report both average accuracy (A_{Mean}) and last accuracy (A_{Last}).

r	λ							
	0.99		0.995		0.999		0.9995	
	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}	A_{Mean}	A_{Last}
0.2	61.30	44.39	61.06	44.05	60.69	43.54	60.26	43.18
0.5	62.12	45.36	61.27	45.05	61.72	44.50	61.36	44.28
1.0	62.60	46.13	63.36	47.64	62.46	46.66	61.79	45.44
2.0	62.51	45.38	60.24	45.80	62.85	46.16	62.50	44.85
5.0	56.70	36.90	57.15	37.97	58.62	38.83	58.08	38.94
CE	59.96	42.39	-	-	-	-	-	-

experimental observations are consistent with our theoretical analysis: smaller r values lead to weaker improvements for older classes and weaker suppression of newer classes, behaving more similarly to cross-entropy, whereas larger r values amplify both the improvement of older classes and the suppression of newer classes.




Figure 6. Forgetting curves with different replay buffer sizes

5.4. Further Analysis

This section briefly discusses the applicability of TAL beyond traditional CIL settings and analyzes its **time efficiency**.

According to the theoretical analysis in Sec. 4, TAL degenerates to the standard cross-entropy loss under two idealized conditions: (1) the training data are perfectly class-balanced; and (2) the samples of each class arrive in a temporally uniform manner. However, in practice, real-world datasets are rarely class-balanced; moreover, even in standard supervised learning without CIL, the randomly shuffled training samples are often not strictly uniform over time. Based on these observations, we further evaluate the effectiveness of TAL under alternative settings, including:

- CIL on long-tailed datasets **without exemplars**, using the **stronger pretrained ViT backbone**.
- comparison between TAL and CE in standard supervised learning **without CIL**.

We also evaluate the time efficiency of TAL, and the results show that its additional computational overhead is minimal (about **0.8%**).




Figure 7. The impact of different r values on the accuracy of each class, arranged in learning order at the end of the last task. Here, $\Delta\text{Accuracy}$ is defined as the difference between the class accuracy of TAL and that of CE (*i.e.*, $\Delta\text{Accuracy} = \text{Accuracy}_{\text{TAL}} - \text{Accuracy}_{\text{CE}}$). It can be observed that TAL does not simply apply uniform protection to all old classes. In the right figure, the more recent old classes (80–89) are even suppressed.

Complete experimental results and detailed analyses are provided in the Appendix.

6. Conclusion

This work revisits the root causes of prediction bias in Class-Incremental Learning (CIL) and argues that, beyond the widely discussed intra-task class imbalance, the *temporal imbalance* between positive and negative supervision plays a crucial role in driving models to suppress early classes and over-favor recently introduced classes. To address this, we propose **Temporal-Adjusted Loss (TAL)**, which performs dynamic temporal modeling for each class through a time-decay memory kernel. Theoretically, TAL degenerates to cross-entropy when sample arrivals are temporally uniform and the data are class-balanced, ensuring stability and compatibility. Empirically, TAL significantly mitigates forgetting and improves overall performance without modifying model architectures, yielding more persistent *feature-level stability* rather than acting solely on the classifier head. Moreover, TAL is **plug-and-play** and can be seamlessly integrated into existing CIL frameworks.

Limitations. While TAL models temporal effects during training, it simplifies the decay of sample influence over time into a fixed exponential form. In reality, the memory parameter λ may vary across tasks or evolve dynamically over time. Additionally, the exponential decay approaches zero as time goes to infinity, which may not fully capture the persistence of learned representations. Future work will explore more flexible, non-parametric forms of temporal modeling to enhance both the interpretability and generality of TAL.

References

- [1] Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [2] Manuel Baena-García, José del Campo-Ávila, Raul Fidalgo, Albert Bifet, Ricard Gavaldà, and Rafael Morales-Bueno. Early drift detection method. In *Fourth international workshop on knowledge discovery from data streams*, pages 77–86, 2006. 2
- [3] Eden Belouadah and Adrian Popescu. I12m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 583–592, 2019. 2
- [4] Albert Bifet and Ricard Gavaldà. Learning from time-changing data with adaptive windowing. In *Proceedings of the 2007 SIAM international conference on data mining*, pages 443–448. SIAM, 2007. 2
- [5] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, Zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, pages 446–461. Springer, 2014. 6
- [6] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 233–248, 2018. 1
- [7] 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. 2
- [8] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer vision—ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part XX 16*, pages 86–102. Springer, 2020. 2
- [9] Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9285–9295, 2022. 2
- [10] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. 1
- [11] Joao Gama, Pedro Medas, Gladys Castillo, and Pedro Rodrigues. Learning with drift detection. In *Brazilian symposium on artificial intelligence*, pages 286–295. Springer, 2004. 2
- [12] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM computing surveys (CSUR)*, 46(4):1–37, 2014. 2
- [13] Qiankun Gao, Chen Zhao, Yifan Sun, Teng Xi, Gang Zhang, Bernard Ghanem, and Jian Zhang. A unified continual learning framework with general parameter-efficient tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11483–11493, 2023. 17
- [14] Rosana Noronha Gemaque, Albert França Josuá Costa, Rafael Giusti, and Eulanda Miranda Dos Santos. An overview of unsupervised drift detection methods. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(6):e1381, 2020. 2
- [15] Jiangpeng He, Xiaoyan Zhang, Luotao Lin, Jack Ma, Heather A. Eicher-Miller, and Fengqing Zhu. Long-tailed continual learning for visual food recognition, 2025. 16, 17
- [16] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8340–8349, 2021. 16, 17
- [17] Charles C Holt. Forecasting seasonals and trends by exponentially weighted moving averages. *International journal of forecasting*, 20(1):5–10, 2004. 2
- [18] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 831–839, 2019. 2
- [19] Rob J Hyndman and George Athanasopoulos. *Forecasting: principles and practice*. OTexts, 2018. 2
- [20] Rob J Hyndman, Anne B Koehler, Ralph D Snyder, and Simone Grose. A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of forecasting*, 18(3):439–454, 2002. 2
- [21] 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. 1, 2
- [22] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, Toronto, Ontario, Canada, 2009. 6
- [23] Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 312–321, 2019. 2
- [24] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 2
- [25] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019. 16
- [26] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, pages 109–165. Elsevier, 1989. 1
- [27] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018. 6

- [28] Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and efficient off-policy reinforcement learning. *Advances in neural information processing systems*, 29, 2016. 2
- [29] 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*, pages 2001–2010, 2017. 1, 2, 6, 7
- [30] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018. 2
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 6
- [32] Harm Seijen and Rich Sutton. True online td (λ). In *International Conference on Machine Learning*, pages 692–700. PMLR, 2014. 2, 4
- [33] James Seale Smith, Leonid Karlinsky, Vyshnavi Gutta, Paola Cascante-Bonilla, Donghyun Kim, Assaf Arbelle, Rameswar Panda, Rogerio Feris, and Zsolt Kira. Coda-prompt: Continual decomposed attention-based prompting for rehearsal-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11909–11919, 2023. 17
- [34] Hai-Long Sun, Da-Wei Zhou, De-Chuan Zhan, and Han-Jia Ye. Pilot: A pre-trained model-based continual learning toolbox, 2025. 16
- [35] Hai-Long Sun, Da-Wei Zhou, Hanbin Zhao, Le Gan, De-Chuan Zhan, and Han-Jia Ye. Mos: Model surgery for pre-trained model-based class-incremental learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 20699–20707, 2025. 17
- [36] Richard S Sutton. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988. 2
- [37] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. MIT press Cambridge, 1998. 2
- [38] 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*, pages 99–108, 2022. 2
- [39] Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 1
- [40] Hado van Hasselt, Sephora Madjiheurem, Matteo Hessel, David Silver, André Barreto, and Diana Borsa. Expected eligibility traces. In *Proceedings of the AAAI conference on artificial intelligence*, pages 9997–10005, 2021. 2
- [41] Joao Vinagre. Time-aware collaborative filtering: a review. In *Doctoral symposium in informatics engineering*, page 43, 2012. 3
- [42] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 16, 17
- [43] Fu-Yun Wang, Da-Wei Zhou, Liu Liu, Han-Jia Ye, Yatao Bian, De-Chuan Zhan, and Peilin Zhao. Beef: Bi-compatible class-incremental learning via energy-based expansion and fusion. In *The eleventh international conference on learning representations*, 2022. 2
- [44] Fu-Yun Wang, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Foster: Feature boosting and compression for class-incremental learning. In *European conference on computer vision*, pages 398–414. Springer, 2022. 2, 6, 7
- [45] Zifeng Wang, Zizhao Zhang, Sayna Ebrahimi, Ruoxi Sun, Han Zhang, Chen-Yu Lee, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, et al. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *European conference on computer vision*, pages 631–648. Springer, 2022. 17
- [46] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 139–149, 2022. 17
- [47] Peter R Winters. Forecasting sales by exponentially weighted moving averages. *Management science*, 6(3):324–342, 1960. 2
- [48] Pei Wu, Chi Ho Yeung, Weiping Liu, Cihang Jin, and Yi-Cheng Zhang. Time-aware collaborative filtering with the piecewise decay function. *arXiv preprint arXiv:1010.3988*, 2010. 3
- [49] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 374–382, 2019. 1, 2
- [50] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3014–3023, 2021. 1, 2, 6, 7
- [51] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13208–13217, 2020.
- [52] Bowen Zheng, Da-Wei Zhou, Han-Jia Ye, and De-Chuan Zhan. Task-agnostic guided feature expansion for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10099–10109, 2025. 2, 6, 7
- [53] Da-Wei Zhou, Qi-Wei Wang, Han-Jia Ye, and De-Chuan Zhan. A model or 603 exemplars: Towards memory-efficient class-incremental learning. In *ICLR*, 2023. 2, 6, 7
- [54] Da-Wei Zhou, Hai-Long Sun, Jingyi Ning, Han-Jia Ye, and De-Chuan Zhan. Continual learning with pre-trained models: A survey. *arXiv preprint arXiv:2401.16386*, 2024. 6
- [55] Da-Wei Zhou, Zi-Wen Cai, Han-Jia Ye, De-Chuan Zhan, and Ziwei Liu. Revisiting class-incremental learning with pre-trained models: Generalizability and adaptivity are all

you need. *International Journal of Computer Vision*, 133(3):
1012–1032, 2025. [2](#)

7. Appendix

7.1. Proof of Theorem 1

We provide a proof of Theorem 1 under the notation used in Sec. 4, where N denotes the total number of steps, $a_k[n] \in \{+1, -1\}$ is the supervision polarity for class k at step n , and $S_k[n]$ denotes the cumulative number of positive samples for class k up to step n . Recall that

$$Q_k[N] = \sum_{n=0}^{N-1} f[N-1-n] a_k[n], \quad (13)$$

where $f[\cdot]$ is a monotonically decreasing memory kernel with $f[n+1] \leq f[n]$.

Step 1: Transforming the supervision polarity. Let $b_k[n] = \frac{a_k[n]+1}{2} \in \{0, 1\}$ so that $a_k[n] = 2b_k[n] - 1$. Substituting into (13) gives

$$Q_k[N] = 2 \sum_{n=0}^{N-1} f[N-1-n] b_k[n] - \sum_{n=0}^{N-1} f[N-1-n]. \quad (14)$$

Step 2: Summation by parts. Define the cumulative positive count

$$S_k[n] := \sum_{j=0}^n b_k[j], \quad n = 0, \dots, N-1,$$

with $S_k[-1] := 0$. We apply summation by parts:

$$\begin{aligned} \sum_{n=0}^{N-1} f[N-1-n] b_k[n] &= f[0] S_k[N-1] - \\ &\sum_{n=0}^{N-2} (f[N-2-n] - f[N-1-n]) S_k[n]. \end{aligned} \quad (15)$$

Step 3: Defining auxiliary statistics. Let

$$\Delta_n := f[N-2-n] - f[N-1-n] \geq 0, \quad n = 0, \dots, N-2, \quad (16)$$

and

$$\Phi_k(N) := f[0] S_k[N-1] - \sum_{n=0}^{N-2} \Delta_n S_k[n]. \quad (17)$$

Then from (14) and (15),

$$Q_k[N] = 2 \Phi_k(N) - \sum_{n=0}^{N-1} f[N-1-n]. \quad (18)$$

Step 4: Order equivalence. The second term in (18) is independent of k . Thus for any two classes A and B ,

$$Q_A[N] \leq Q_B[N] \iff \Phi_A(N) \leq \Phi_B(N). \quad (19)$$

Step 5: Equal positive counts. Let $M_k := S_k[N-1]$ be the total number of positive samples for class k . If $M_A = M_B = M$, (17) and (19) yield

$$Q_A[N] \leq Q_B[N] \iff \sum_{n=0}^{N-2} \Delta_n S_A[n] \geq \sum_{n=0}^{N-2} \Delta_n S_B[n]. \quad (20)$$

Step 6: Back-loading condition. Under the back-loading condition

$$S_A[n] \geq S_B[n], \quad \forall n = 0, \dots, N-2,$$

and since $\Delta_n \geq 0$ by the monotonicity of $f[\cdot]$, it follows that

$$\sum_{n=0}^{N-2} \Delta_n S_A[n] \geq \sum_{n=0}^{N-2} \Delta_n S_B[n],$$

and therefore $Q_A[N] \leq Q_B[N]$ by (20).

Step 7: Strict inequality. If $f[\cdot]$ is strictly decreasing, then $\Delta_n > 0$ for all n , and strict inequality holds whenever $\exists n$ such that $S_A[n] > S_B[n]$.

Conclusion. This completes the proof of Theorem 1, showing that for two classes with the same total number of positive samples, the class with later (back-loaded) positives obtains a larger temporal supervision strength $Q_k[N]$.

7.2. Invariant Range and Boundary Attainment of Q and $w(\cdot)$

We show that the interval $[0, Q_{\max})$ is invariant for the recursion in Eq. 10 and that $w(Q_k[N]) \in [0, 1)$ is preserved as well, together with the extremal sequences that attain the lower/upper boundary (in the appropriate sense).

Setup. Recall the recursion (Eq. 10)

$$Q_k[N+1] = \lambda \left(Q_k[N] + \begin{cases} +1, & a_k[N] = +1, \\ -w(Q_k[N]), & a_k[N] = -1, \end{cases} \right), \quad (21)$$

with $0 < \lambda < 1$, $Q_k[0] = 0$, $Q_{\max} = \frac{\lambda}{1-\lambda}$ and

$$w(q) = (q/Q_{\max})^r, \quad r > 0.$$

Throughout we assume $r \geq 1$ and $Q_{\max} \geq 1$ (equivalently, $\lambda \geq \frac{1}{2}$).

Proposition 1 (Invariance of $[0, Q_{\max})$ and $[0, 1)$). *If $0 \leq Q_k[N] < Q_{\max}$ and $0 \leq w(Q_k[N]) < 1$, then*

$$0 \leq Q_k[N+1] < Q_{\max}, \quad 0 \leq w(Q_k[N+1]) < 1.$$

Proof. *Upper bound:* Assume $Q_k[N] < Q_{\max}$.

(i) If $a_k[N] = +1$,

$$\begin{aligned} Q_k[N+1] &= \lambda(Q_k[N] + 1) \\ &\leq \lambda(Q_{\max} + 1) \\ &= \lambda\left(\frac{\lambda}{1-\lambda} + 1\right) \\ &= Q_{\max}, \end{aligned} \quad (22)$$

and the inequality is strict since $\lambda < 1$ and $Q_k[N] < Q_{\max}$.

(ii) If $a_k[N] = -1$, since $0 \leq w(Q_k[N]) < 1$,

$$\begin{aligned} Q_k[N+1] &= \lambda(Q_k[N] - w(Q_k[N])) \\ &\leq \lambda Q_k[N] \\ &< \lambda Q_{\max} \\ &< Q_{\max}. \end{aligned} \quad (23)$$

Lower bound: Let $q := Q_k[N] \in [0, Q_{\max})$. With $r \geq 1$ and $Q_{\max} \geq 1$,

$$w(q) = \left(\frac{q}{Q_{\max}}\right)^r \leq \frac{q}{Q_{\max}} \leq q, \quad (24)$$

so $q - w(q) \geq 0$.

(i) If $a_k[N] = +1$,

$$Q_k[N+1] = \lambda(q+1) \geq 0.$$

(ii) If $a_k[N] = -1$,

$$Q_k[N+1] = \lambda(q - w(q)) \geq 0.$$

Thus, $0 \leq Q_k[N+1] < Q_{\max}$. Finally, since $Q_k[N+1]/Q_{\max} \in [0, 1)$ and $r > 0$,

$$0 \leq w(Q_k[N+1]) = \left(\frac{Q_k[N+1]}{Q_{\max}}\right)^r < 1. \quad (25)$$

□

Boundary attainment. *Lower boundary (attained).* If $a_k[n] \equiv -1$ for all $n \geq 0$, then

$$\begin{aligned} Q_k[1] &= \lambda(0 - 0) = 0, \\ Q_k[2] &= \lambda(0 - 0) = 0, \\ &\vdots \end{aligned} \quad (26)$$

so $Q_k[N] \equiv 0$ for all N , and $w(Q_k[N]) \equiv 0$.

Upper boundary (approached, not attained). If $a_k[n] \equiv +1$ for all $n \geq 0$,

$$Q_k[N+1] = \lambda(Q_k[N] + 1), \quad Q_k[0] = 0. \quad (27)$$

The closed-form solution is

$$\begin{aligned} Q_k[N] &= \lambda \sum_{m=0}^{N-1} \lambda^m \\ &= \frac{\lambda(1 - \lambda^N)}{1 - \lambda} \\ &= Q_{\max}(1 - \lambda^N). \end{aligned} \quad (28)$$

Hence $Q_k[N] \neq Q_{\max}$ for finite N , but

$$Q_k[N] \uparrow Q_{\max} \quad \text{as } N \rightarrow \infty.$$

Correspondingly,

$$w(Q_k[N]) = \left(1 - \lambda^N\right)^r \xrightarrow{N \rightarrow \infty} 1^-. \quad (29)$$

Conclusion. Under the stated conditions, the recursion preserves

$$Q_k[N] \in [0, Q_{\max}), \quad w(Q_k[N]) \in [0, 1),$$

for all N . The all-negative sequence attains the lower boundary exactly, while the all-positive sequence approaches the upper boundary asymptotically and never reaches it for $0 < \lambda < 1$.

7.3. Calibration Proof and Closed-Form α for $r \in \{1, 2\}$

Recall the steady-state calibration under temporally uniform and class-balanced data with prior $p = 1/C$:

$$\alpha \cdot w(Q^*) = 1, \quad w(q) = \left(\frac{q}{Q_{\max}}\right)^r,$$

and the steady-state recursion (taking expectation in Eq. 10 and imposing $Q_k[N+1] = Q_k[N] = Q^*$):

$$Q^* = \lambda(Q^* + p - (1-p)w(Q^*)). \quad (30)$$

Let $x^* := Q^*/Q_{\max} \in (0, 1)$ with $Q_{\max} = \lambda/(1-\lambda)$. Dividing (30) by Q_{\max} yields

$$(1-p)(x^*)^r + x^* - p = 0. \quad (31)$$

For $p = 1/C$ this becomes

$$\left(1 - \frac{1}{C}\right)(x^*)^r + x^* - \frac{1}{C} = 0, \quad \alpha = \frac{1}{(x^*)^r}. \quad (32)$$

Existence and uniqueness of $x^* \in (0, 1)$. Define $g(x) = (1-p)x^r + x - p$ on $[0, 1]$. Then $g(0) = -p < 0$, $g(1) = 2(1-p) > 0$ (for $C \geq 2$), and $g'(x) = (1-p)rx^{r-1} + 1 > 0$. Hence g is strictly increasing and crosses zero exactly once on $(0, 1)$; therefore x^* and $\alpha = 1/(x^*)^r$ are unique.

Closed-form for $r = 1$. When $r = 1$, (32) is linear:

$$\left(1 - \frac{1}{C}\right)x^* + x^* - \frac{1}{C} = 0 \implies x^* = \frac{1/C}{2 - 1/C} = \frac{1}{2C - 1}.$$

Hence

$$\alpha = \frac{1}{x^*} = 2C - 1.$$

Closed-form for $r = 2$. When $r = 2$, (32) is quadratic:

$$\left(1 - \frac{1}{C}\right)(x^*)^2 + x^* - \frac{1}{C} = 0.$$

The root in $(0, 1)$ is

$$x^* = \frac{-1 + \sqrt{1 + 4\frac{1}{C}\left(1 - \frac{1}{C}\right)}}{2\left(1 - \frac{1}{C}\right)} = \frac{-C + \sqrt{C^2 + 4C - 4}}{2(C - 1)}.$$

Therefore

$$\alpha = \frac{1}{(x^*)^2} = \left(\frac{C + \sqrt{C^2 + 4C - 4}}{2}\right)^2.$$

Other r . For general $r > 4$, (32) does not admit an elementary closed form in general. Since the left-hand side is strictly increasing on $[0, 1]$ and changes sign between 0 and 1, one can solve for x^* robustly by bracketing (e.g., bisection) or use Newton’s method

$$x^{(t+1)} = x^{(t)} - \frac{\left(1 - \frac{1}{C}\right)(x^{(t)})^r + x^{(t)} - \frac{1}{C}}{\left(1 - \frac{1}{C}\right)r(x^{(t)})^{r-1} + 1},$$

initialized, e.g., at $x^{(0)} = \frac{1}{C}$. The unique solution $x^* \in (0, 1)$ then yields $\alpha = 1/(x^*)^r$.

7.4. Pseudo Code of TAL

In this section, we summarize the batched implementation of the proposed **Temporal-Adjusted Loss (TAL)** in a clean and structured manner. Given a minibatch with logits $Z \in \mathbb{R}^{N \times C}$ and labels $y \in \{1, \dots, C\}^N$, TAL dynamically adjusts the supervision strength for negative samples based on the temporal supervision vector \mathbf{Q} . Specifically, the loss reweights *only* the negative logits using the multiplicative factor $\alpha w(Q_k)$, where α is determined by the number of classes C and the temporal adjustment parameter r .

\mathbf{Q} is initialized as a zero vector at the beginning of training and is updated online after every minibatch to track the relative temporal frequency of each class. The algorithm 1 below illustrates the forward loss computation and the update of \mathbf{Q} in each training step.

7.5. Time Efficiency

7.5.1. Theoretical Complexity Analysis

The proposed Temporal-Adjusted Loss (TAL) introduces two additional computations on top of the standard Cross-Entropy (CE) loss: (i) a temporal reweighting of negative logits via elementwise vector operations, and (ii) an on-line update of the temporal supervision vector \mathbf{Q} after each minibatch. Let N denote the batch size and C the number of classes.

Cross-Entropy (CE). The standard CE loss involves computing the log-sum-exp across all C classes for each of the N samples:

$$\mathcal{O}_{\text{CE}} = \mathcal{O}(NC),$$

dominated by the log-sum-exp and gather operations.

TAL. Compared to CE, TAL adds: (1) a single vectorized elementwise operation to compute $\mathbf{s} = (Q/Q_{\max})^r$ and $\ell = \log(\alpha \mathbf{s})$, and (2) a vector update of \mathbf{Q} using minibatch counts. Both steps are $\mathcal{O}(C)$ operations, which are negligible compared to $\mathcal{O}(NC)$. **Thus, the total complexity is:**

$$\mathcal{O}_{\text{TAL}} = \mathcal{O}(NC) + \mathcal{O}(C) = \mathcal{O}(NC),$$

identical in order to CE.

7.5.2. Total Training Time Evaluation

To assess the actual runtime impact of TAL in a realistic setting, we measure the *total wall-clock training time* of several representative CIL methods on CIFAR-100, with and without TAL (i.e., using standard CE loss). This includes all forward and backward passes, model expansions, memory operations, and optimizer updates across the full training process.

We use the following setting:

- **Dataset:** CIFAR-100
- **Number of classes per task:** 10
- **Exemplars per class:** 20
- **Batch size:** 128
- **Backbone:** ResNet-32
- **GPU:** NVIDIA A40, 46 GB

The results in Table 3 show that integrating TAL into existing methods introduces only marginal runtime overhead. On average, the total training time increases by **0.76%**. This confirms that TAL’s lightweight computation has minimal impact at the full training scale.

Observation. TAL adds less than 1% overhead on average over a complete training run, demonstrating that its temporal supervision mechanism does not slow down practical CIL pipelines. Meanwhile, Meanwhile, we also observe that BiC, which relies on balanced finetuning, significantly increases the total training time and computational

Algorithm 1 TAL’s Pseudo-code (Vectorized Form)

Require: Logits $Z \in \mathbb{R}^{N \times C}$, targets $y \in \{1, \dots, C\}^N$, temporal supervision vector $\mathbf{Q} = (Q_1, \dots, Q_C)^\top \in \mathbb{R}^C$, hyperparameters $\lambda \in (0, 1)$, $r > 0$, α determined by C and r

Initialization:

- 1: $\mathbf{Q} \leftarrow \mathbf{0}$ ▷ Initialize \mathbf{Q} as a zero vector at the beginning of training
- 2: $Q_{\max} \leftarrow \frac{\lambda}{1 - \lambda}$

Part A: Loss computation

- 3: $s_k \leftarrow \left(\frac{Q_k}{Q_{\max}}\right)^r$ for $k = 1, \dots, C$ ▷ $s_k \in [0, 1]$
- 4: $\ell_k \leftarrow \log(\alpha \cdot \max(s_k, \varepsilon))$ for $k = 1, \dots, C$ ▷ stabilized log-weight
- 5: $\tilde{Z}_{i,k} \leftarrow Z_{i,k} + \ell_k$ for all $i = 1, \dots, N$ and $k = 1, \dots, C$
- 6: $\tilde{Z}_{i,y_i} \leftarrow Z_{i,y_i}$ for all $i = 1, \dots, N$ ▷ vectorized replacement for true classes
- 7: $\mathcal{L}_{\text{TAL}} \leftarrow \frac{1}{N} \sum_{i=1}^N \left[\log \left(\sum_{k=1}^C e^{\tilde{Z}_{i,k}} \right) - Z_{i,y_i} \right]$

Part B: Update of \mathbf{Q}

- 8: $N_p(k) \leftarrow \sum_{i=1}^N \mathbb{1}\{y_i = k\}$ for $k = 1, \dots, C$
 - 9: $N_n(k) \leftarrow N - N_p(k)$ for $k = 1, \dots, C$
 - 10: *Vector form:* $\mathbf{Q} \leftarrow \lambda \left(\mathbf{Q} + \frac{\mathbf{N}_p}{N} - \frac{\mathbf{N}_n}{N} \cdot \mathbf{s} \right)$
 - 11: *Component-wise (for clarity):* $Q_k \leftarrow \lambda \left(Q_k + \frac{N_p(k)}{N} - \frac{N_n(k)}{N} \cdot s_k \right)$, $k = 1, \dots, C$
 - 12: **return** $\mathcal{L}_{\text{TAL}}, \mathbf{Q}$
-

Table 3. Comparison of **total training time** on CIFAR-100 between Cross-Entropy (CE) and TAL. The third column reports the relative change (%) in total training time.

Method	CE (min)	TAL (min)	% Δ (TAL vs CE)
ER	42.5	42.8	+0.7%
BiC	208.9	211.7	+1.3%
iCaRL	44.2	45.2	+2.3%
FOSTER	66.7	67.8	+1.6%
MEMO	79.6	78.1	-1.9%
DER	81.3	82.1	+1.0%
TagFex	125.1	125.1	0.0%
Average	92.6	93.3	+0.76%

cost. This highlights the advantage of TAL in correcting bias without requiring any additional post-processing.

7.5.3. Loss-Only Runtime Evaluation

While the previous experiment evaluates full training time, we further isolate the *loss computation step* to examine TAL’s intrinsic overhead compared to CE. Specifically, we benchmark the per-batch time of CE and TAL loss computation under different batch sizes and numbers of classes, *excluding forward propagation and any CIL-specific com-*

ponents such as memory replay or model updates. This allows us to measure the pure cost of the loss function itself.

As shown in Table 4, TAL is slightly slower than CE due to the additional $\mathcal{O}(C)$ vector operations, but the gap remains constant and negligible across both batch size and class count. This empirically verifies the theoretical complexity analysis in Section 7.5.

Table 4. **Per-batch loss computation time** comparison between Cross-Entropy (CE) and TAL under different batch sizes and numbers of classes. Units in 10^{-2} ms.

Batch	Method	5 cls	20 cls	100 cls	500 cls
32	CE	10.89	10.78	10.77	10.81
	TAL	42.77	42.87	42.77	43.03
64	CE	10.83	10.86	10.76	10.81
	TAL	42.90	42.85	43.13	43.04
128	CE	10.72	10.74	10.75	10.79
	TAL	42.89	42.80	42.91	43.63
256	CE	10.80	10.81	10.82	11.08
	TAL	42.90	42.81	43.30	43.85

Observation. Even in this isolated setting, TAL remains

efficient. The per-batch runtime difference does not scale with batch size or number of classes, confirming its constant and minimal overhead.

7.6. TAL with Pretrained Model Based CIL Methods

This section examines the performance of **TAL when applied to Pretrained Model Based CIL methods (PTM methods)**.

In the main experiments of this paper, we assumed no pretrained backbone and trained models from scratch. In contrast, PTM methods rely on a *pretrained vision transformer backbone*, where the feature extractor is pretrained and frozen throughout continual learning, eliminating feature drift between tasks.

Moreover, PTM methods are typically *exemplar-free*. This implies that:

- Each task is treated as an independent classification problem.
- Old classes are excluded from the loss computation during subsequent tasks.
- Each class receives supervision only in its original task and no further positive or negative supervision thereafter.

As a result, the *positive-negative supervision imbalance* analyzed in Sec. 4 does not apply under the PTM setting.

However, as discussed in Sec. 6, TAL reduces to CE only if two conditions are met: (1) the training samples of each class are temporally uniform, and (2) the dataset is class-balanced. While condition (1) is satisfied in PTM scenarios, condition (2) often is not—meaning that TAL remains effective for class-imbalanced datasets.

Tab. 5 reports results on three long-tailed datasets under 5-task and 10-task settings. We observe consistent improvements in both overall accuracy and tail class accuracy across multiple PTM-based baselines, demonstrating the general applicability of TAL even in pretrained settings without exemplars. A_{Mean} denotes the average accuracy over all classes, while A_{Tail} measures the accuracy on the tail classes, defined as the bottom 50% of classes ranked by the number of training samples. All numbers are averaged on 3 different random seeds. Same as the Sec. 5.2, the TAL hyperparameters are fixed as $\lambda = 0.995$ and $r = 1$.

Implementation Details. For experiments with PTM methods, as shown in 8, we adopt three long-tailed datasets: **CUB-LT** [42], **ImageNetR-LT** [16], and **Food101-LT** [15]. Each long-tailed training set is constructed by resampling the original dataset following a Pareto distribution with shape parameter $\alpha = 6$, as in [25]. The imbalance factors $\rho = n_{\text{max}}/n_{\text{min}}$ are set to 20 for CUB-LT, 81 for ImageNetR-LT, and 150 for Food101-LT. The test sets are class-balanced, containing 20 samples per class for

CUB-LT and 25 samples per class for ImageNetR-LT and Food101-LT.

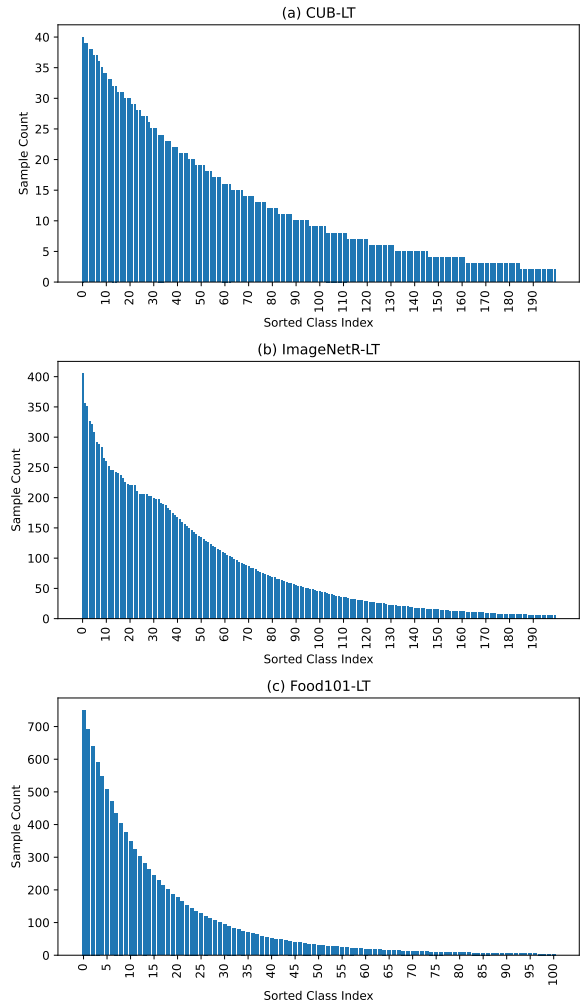


Figure 8. Sample-Class Distribution of Long-Tailed Datasets

All PTM experiments use a **frozen pretrained ViT backbone**, following standard pretrained model-based CIL settings. No exemplars are used, and each task introduces a disjoint set of classes without overlap. We adopt the default hyperparameter configurations provided in the **LAMDA-PILOT** [34] framework for all methods, including learning rate, optimizer settings, and training schedule. The number of tuning epochs for each task is fixed at 20 across all PTM methods to ensure a fair comparison.

Observations Although PTM settings do not involve positive-negative supervision imbalance for the same class across different tasks, tail classes within each task still suffer from excessively strong negative supervision. TAL addresses this issue without relying on any prior class knowledge or post-processing, and consistently improves both

Table 5. Average accuracy (A_{Mean}) and tail class accuracy (A_{Tail}) under 5-task and 10-task settings across three long-tailed datasets. +TAL implies that cross entropy loss is replaced with TAL in the PTM-based methods.

Method	CUB-LT [42]				ImageNetR-LT [16]				Food101-LT [15]			
	5-task		10-task		5-task		10-task		5-task		10-task	
	A_{Mean}	A_{Tail}	A_{Mean}	A_{Tail}	A_{Mean}	A_{Tail}	A_{Mean}	A_{Tail}	A_{Mean}	A_{Tail}	A_{Mean}	A_{Tail}
L2P (CVPR'22)[46]	74.22	57.20	68.91	46.05	70.03	56.32	69.94	54.52	55.36	29.33	47.46	16.11
+ TAL	75.91	61.19	70.71	49.80	72.84	63.50	71.63	60.47	58.53	40.93	50.84	23.84
DualPrompt (ECCV'22)[45]	73.85	51.30	69.87	50.75	65.40	48.28	65.14	46.52	52.24	25.55	49.84	22.56
+ TAL	75.06	54.65	70.68	52.50	66.45	52.04	66.00	50.06	54.28	30.01	50.77	24.89
CODAPrompt (CVPR'23)[33]	74.14	53.26	69.10	49.54	72.92	57.16	70.61	53.60	62.64	34.71	56.11	28.33
+ TAL	75.51	56.04	70.45	51.68	75.34	62.08	72.86	58.00	65.19	39.44	58.12	34.10
LAE (ICCV'23) [13]	74.54	57.90	72.03	50.20	69.13	53.52	69.54	56.00	62.14	37.96	56.43	29.91
+ TAL	76.97	61.40	74.83	57.10	70.54	58.48	71.00	58.64	64.34	42.35	59.38	35.89
MOS (AAAI'25)[35]	73.95	51.40	71.52	48.93	69.15	53.80	68.73	55.44	60.63	29.30	55.46	28.08
+ TAL	75.30	54.86	72.14	50.95	71.47	59.29	70.31	58.77	63.24	33.27	57.92	33.73

A_{Mean} and A_{Tail} on long-tailed datasets. This demonstrates that TAL is effective beyond non-PTM scenarios and can be seamlessly integrated into pretrained pipelines without architectural or training modifications.

7.7. TAL in Standard Supervised Learning

As discussed in Sec. 6, TAL reduces to CE only if two conditions are met: (1) the training samples of each class are temporally uniform, and (2) the dataset is class-balanced. This implies that if both conditions hold—as in standard supervised learning settings where samples are randomly shuffled and the class distribution is balanced—TAL should theoretically behave similarly to CE.

However, we consistently observe in practice that even on the first task, where no forgetting can occur, TAL often yields slightly higher test accuracy than CE when appropriately tuned. This phenomenon caught our attention, as the first task represents a purely supervised learning scenario without any temporal imbalance. To verify this effect, we conducted experiments on CIFAR-100 and ImageNet, where we randomly selected subsets of classes and trained models using ResNet-32 and ResNet-18 backbones from scratch, respectively. We used identical optimizers and training schedules for both CE and TAL to ensure a fair comparison.

Tab. 6 summarizes the results under different batch sizes and class numbers. Across all configurations, TAL consistently outperforms CE by a small but stable margin. This suggests that TAL not only maintains theoretical equivalence to CE under balanced conditions but may also provide a mild regularization effect through its temporal weighting mechanism, leading to slightly improved generalization. We believe this observation opens up interesting directions for applying TAL beyond continual learning. One possible explanation for this phenomenon is the existence of a subtle

Table 6. Average accuracy (%) comparison between Cross-Entropy (CE) and TAL under different batch sizes and class numbers on CIFAR and ImageNet. **Classes are randomly chosen with 10 different random seeds.**

Dataset	Method	Batch=64	Batch=128	Batch=256
5 classes				
CIFAR	CE	88.56	88.68	85.88
	TAL	89.36	89.24	89.36
ImageNet	CE	96.83	94.84	96.43
	TAL	97.62	95.24	97.22
10 classes				
CIFAR	CE	85.64	84.56	84.28
	TAL	87.42	85.76	84.88
ImageNet	CE	89.00	86.60	88.24
	TAL	89.80	88.40	88.68

inter-batch forgetting effect. Specifically, since each epoch in standard supervised learning involves random shuffling of the training data, some classes may have a larger proportion of their samples appearing in earlier batches of an epoch. As a result, by the end of the epoch, these classes may have received relatively stronger negative supervision, leading to a mild temporal imbalance effect that resembles, to a lesser degree, the temporal imbalance observed in continual learning scenarios.

This finding suggests that TAL’s temporal adjustment mechanism can have beneficial effects even in conventional supervised learning, by mitigating these subtle temporal biases during optimization.

7.8. TAL with Fixed Buffer Size

Table 7 reports the results on ImageNet-100 under two class-incremental settings. Here, $Bx\text{ Inc}y$ denotes that the base task contains x classes and each incremental task introduces y new classes. All methods use a fixed replay buffer with 2000 exemplars in total, and “+ Ours” indicates the baseline method with the Temporal-Adjusted Loss applied.

Table 7. Average accuracy (A_{Mean}) and last accuracy (A_{Last}) on ImageNet-100 under the B10 Inc10 setting. Each entry is reported as *baseline / +Ours*.

Method	A_{Mean}	A_{Last}
iCaRL	63.06 / 71.29	51.37 / 55.27
FOSTER	69.81 / 73.27	56.72 / 60.02
MEMO	67.83 / 73.65	57.45 / 60.64
DER	75.38 / 78.73	64.49 / 67.17
TagFex	78.95 / 81.32	66.37 / 69.72