

Class-incremental Learning with Rectified Feature-Graph Preservation

Cheng-Hsun Lei*, Yi-Hsin Chen*, Wen-Hsiao Peng, and Wei-Chen Chiu

National Chiao Tung University, Taiwan

{raygoah.cs07g, yhchen.iie07g}@nctu.edu.tw {wpeng, walon}@cs.nctu.edu.tw

Abstract. In this paper, we address the problem of distillation-based class-incremental learning with a single head. A central theme of this task is to learn new classes that arrive in sequential phases over time while keeping the model’s capability of recognizing seen classes with only limited memory for preserving seen data samples. Many regularization strategies have been proposed to mitigate the phenomenon of catastrophic forgetting. To understand better the essence of these regularizations, we introduce a feature-graph preservation perspective. Insights into their merits and faults motivate our weighted-Euclidean regularization for old knowledge preservation. We further propose rectified cosine normalization and show how it can work with binary cross-entropy to increase class separation for effective learning of new classes. Experimental results on both CIFAR-100 and ImageNet datasets demonstrate that our method outperforms the state-of-the-art approaches in reducing classification error, easing catastrophic forgetting, and encouraging evenly balanced accuracy over different classes. Our project page is at : <https://github.com/yhchen12101/FGP-ICL>.

1 Introduction

Class-incremental learning [1] is a difficult yet practical problem for visual recognition models. The task requires the model with a single classification head to learn new classes that arrive sequentially while preserving its knowledge on the old ones. It is worth noting that class-incremental learning is different from task-incremental learning [2,3,4,5,6], which learns to handle multiple tasks (i.e. groups of classes, possibly from different datasets) given sequentially. In particular, task-incremental learning has a unique assumption that the task indices of classes are known during both training and test time. Thus, the model typically aims to learn a shared feature extractor across tasks under the multi-head setup (i.e. each task has its own head). Even though there are some works [3,7,8] addressing class-incremental learning from the perspective of task-incremental learning, the differences stemmed from the aforementioned assumption clearly distinguish these works from typical class-incremental learning. This paper focuses on the more difficult yet practical setup to learn a single-head classifier without task identification.

* Both authors contributed equally to the paper

Generally, class-incremental learning is faced with two major challenges: (1) the model needs to be updated on-the-fly upon the arrival of training data of new classes, and (2) there is only limited memory space (also known as experience reply buffer) for retaining partially the training data of previously learned classes. It is shown in [1] that naively fine-tuning the model on newly received training data of incoming classes without considering the previously learned classes often suffers from *catastrophic forgetting* [9]; that is, the classification accuracy on old classes may decline dramatically when the model attempts to learn new classes.

Several prior works based on deep learning models have been proposed to deal with catastrophic forgetting. They can roughly be categorized into three types: the parameter-based, distillation-based, and generative-model-based approaches. The parameter-based methods aim to restrict the update on network parameters that are highly sensitive to the recognition accuracy of old classes [2,3,10,11], and to conserve the usage of network parameters in learning new classes [12,13,14]. Identifying and regularizing those important network parameters can be prohibitively expensive, especially with large learning models. By contrast, the distillation-based methods leverage the idea of knowledge distillation [15], requiring that the model updated with the training data of new classes should produce similar predictions to the old model over the previously learned classes [1,16,17,18,19,4,20,21]. They typically need a divergence measure to be defined that measures the similarity between the old and new predictions. In particular, this class of methods is often influenced by the amount of data available in the experience reply buffer on which the similarity measurement is conducted. The generative-model-based methods replace the experience replay buffer with deep generative models [22], so that these generative models can later be utilized to generate synthetic training data for old classes while updating the model to learn new classes [23,24,25]. Obviously, their capacity for learning the distribution of old data crucially affects the quality of generated training data.

This work presents a new geometry-informed attempt at knowledge distillation for class-incremental learning. Instead of devising new divergences for similarity measurement, our work stresses preserving the geometric structure of the feature space. Specifically, we first (1) analyze the objectives used in the existing distillation-based algorithms from a *feature-graph preservation* perspective. By inspection of their merits and faults, we then (2) propose *weighted-Euclidean regularization* to preserve the knowledge of previously learned classes while providing enough flexibility for the model to learn new classes. Lastly, we (3) introduce *rectified cosine normalization* to learn discriminative feature representations, together with the objective of binary cross-entropy.

2 Related Work

Parameter-based. Parameter-based methods attempt to identify the network parameters that have a profound impact on the classification accuracy of previously learned classes once altered. These important parameters are encouraged to be fixed during the incremental learning in order to resolve the issue of

catastrophic forgetting (e.g. MAS [2], EWC [3], SI [10], and RWalk [11]). Other works such as Rahaf *et al.* [12], Piggy-back [13], and Packnet [14] further advance freeing up the remaining parameters selectively in learning new classes. This is achieved by learning a weighting mask, pruning the network, or adding specifically-designed regularization. Because separate parts of the network can be activated for different tasks, parameter-based methods often find application in task-incremental learning. They however incur heavy computation in identifying and regularizing important network parameters.

Generative-model-based. Generative-model-based methods, such as [23,24,25], turn to generative models as a means of keeping the training data of old classes. Since the experience replay buffer has a limited size, the training data of previously learned classes can only be retrained partially. This leads to an imbalanced distribution of training data between the new and old classes. To overcome this issue, generative-model-based methods employ generative models to generate synthetic training data for old classes in the course of incremental learning. It is worth noting that the use, storage, and training of these generative models can cause a considerable complexity increase.

Distillation-based. Knowledge distillation [15] is originally proposed to distill the knowledge from a teacher network (which is typically a network with larger capacity) and transfer it to a student network (which is a smaller network than the teacher), by encouraging the student to have a similar posterior distribution to the teacher’s. This concept is adapted to preserve the model’s knowledge of old classes for class-incremental learning. LwF [4] is the first work to apply knowledge distillation to incremental learning, introducing a modified cross-entropy loss function. **iCaRL** [1] uses binary cross-entropy as a new distillation loss, provides a mechanism for selecting old-class exemplars, and classifies images with the *nearest-mean-of-exemplars* criterion. Castro *et al.* [17] proposes a framework, termed **End-to-End** in this paper, that adopts the ordinary cross-entropy as the distillation loss in learning end-to-end the feature representation and the classifier. In particular, it introduces balanced fine-tuning as an extra training step to deal with the data imbalance between the old and new classes. **BIC** [20] proposes a bias correction method to address the data imbalance, achieving excellent accuracy on large datasets. **Hou19** [18] proposes several designs, including the cosine normalization, the less-forget constraint, and the inter-class separation, to encourage the classifier to treat old and new classes more uniformly. Based on **Hou19**’s training strategies, Liu *et al.* [21], termed **Mnemonics**, further optimize the parameterized exemplars by a bi-level optimization program. Both **Hou19** and **Mnemonics** rely on a well-pretrained feature extractor to obtain superior performance on incremental learning.

3 Proposed Method

This section presents our strategies – namely, weighted-Euclidean regularization and rectified cosine normalization – for dealing with catastrophic forgetting and effective new class learning in the scenario of class-incremental learning.

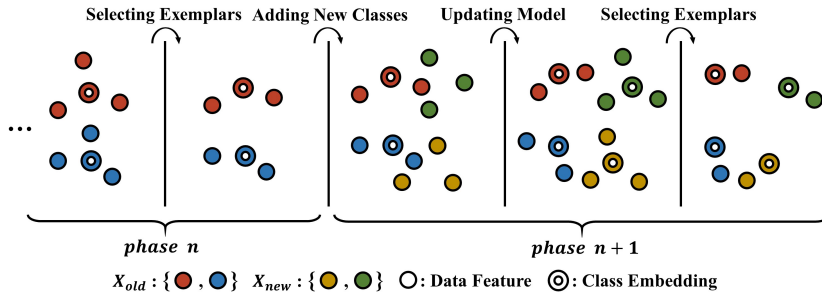


Fig. 1. Illustration of the class-incremental learning process: (1) selecting exemplars, (2) adding the data samples of new classes, and (3) updating the model.

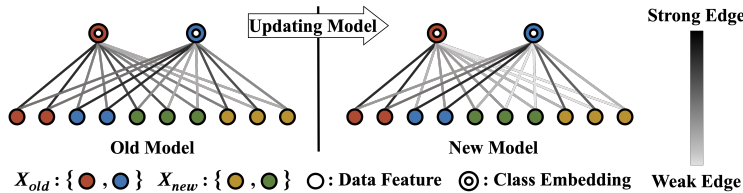


Fig. 2. Knowledge distillation from a bipartite feature-graph preservation perspective, with the bottom row and top row denoting the feature vectors of data samples and the class embeddings respectively. Minimizing the distillation loss is viewed as the structure (i.e. edge strength) preservation of such a bipartite graph during the model update.

Fig. 1 illustrates the learning of data features and class embeddings (i.e. the classifier’s weight vectors) from one incremental learning phase n to the next $n + 1$. As shown, the data of different classes arrive sequentially in different *phases*. We denote the old data that have arrived before phase $n + 1$ as X_{old} and the new data arriving in phase $n + 1$ as X_{new} . Assuming we already have a model capable of classifying the classes C_{old} in the X_{old} , we now aim to train the model to additionally learn the new classes C_{new} in X_{new} . However, there is limited memory for storing data samples; we can only preserve some exemplars chosen from the seen samples before phase $n + 1$. Thus, the training is based on the dataset $X_{all} = X_{new} \cup X_{exem}$, where X_{exem} is composed of the exemplars of C_{old} (i.e X_{exem} is a tiny subset of X_{old}) and the number of exemplars that can be kept is constrained by the memory size.

3.1 Mitigating Catastrophic Forgetting

To motivate the need for our weighted-Euclidean regularization, we introduce a feature-graph preservation perspective on knowledge distillation. This new perspective reveals the essence of existing distillation-based regularization schemes in how they tackle catastrophic forgetting and highlights the striking feature of our method.

Knowledge Distillation in Class-Incremental Learning. We begin by revisiting the prior works from iCaRL [1] and End-to-End [17]. To prevent class-incremental learning from catastrophic forgetting, both models introduce a distillation loss that requires the model after the gradient update (termed the *new* model) to produce a similar prediction on C_{old} to that made by the model before the update (termed the *old* model), for any data in X_{all} . In symbols, for a data sample $x_k \in X_{all}$, iCaRL [1] computes the distillation loss $\mathcal{L}_{dist.bce}(x_k)$ to be the sum of the binary cross-entropy between the old prediction $p_{i|k}^*$ and the new prediction $p_{i|k}$ over different classes $i \in C_{old}$:

$$\mathcal{L}_{dist.bce}(x_k) = \sum_{i \in C_{old}} \left\{ 1 \times \left[-p_{i|k}^* \log p_{i|k} - (1 - p_{i|k}^*) \log(1 - p_{i|k}) \right] \right\}, \quad (1)$$

where $p_{i|k} = \sigma(a_{i|k})$ is the sigmoid output of the activation $a_{i|k} = w_i^T f_k + b_i$ evaluated based on the new model, with w_i, b_i representing the weight vector and bias of the i -th classifier respectively, and f_k denoting the feature vector of x_k ; and $p_{i|k}^* = \sigma(a_{i|k}^*)$ is computed similarly yet with the activation $a_{i|k}^* = w_i^{*T} f_k^* + b_i^*$ evaluated with the old model. In contrast, the Eq. (3) in the End-to-End [17] paper minimizes the KL-divergence¹ between $p_{i|k}^*$ and $p_{i|k}$ over $i \in C_{old}$ for preserving knowledge learned in previous phases:

$$\mathcal{L}_{dist.KL}(x_k) = \sum_{i \in C_{old}} \left\{ p_{i|k}^* \times \left[\log p_{i|k}^* - \log p_{i|k} \right] \right\}, \quad (2)$$

where $p_{i|k} = \text{softmax}(a_{i|k})$ (respectively $p_{i|k}^* = \text{softmax}(a_{i|k}^*)$) is the softmax output of the activation $a_{i|k}/T$ (respectively, $a_{i|k}^*/T$) attenuated by a temperature parameter T .

Close examination of Eqs. (1) and (2) above suggests a more general distillation loss of the form

$$\mathcal{L}_{dist}(x_k) = \sum_{i \in C_{old}} \left\{ \gamma_{i|k} \times \mathcal{D}(p_{i|k}^*, p_{i|k}) \right\}, \quad (3)$$

with $\mathcal{D}(p_{i|k}^*, p_{i|k})$ measuring the discrepancy between the old and new predictions (i.e. $p_{i|k}^*$ and $p_{i|k}$) for class i , and $\gamma_{i|k}$ weighting the contribution of $\mathcal{D}(p_{i|k}^*, p_{i|k})$ to the resulting loss $\mathcal{L}_{dist}(x_k)$. Note that a non-zero discrepancy value $\mathcal{D}(p_{i|k}^*, p_{i|k})$, be it positive or negative, signals a change in the model’s prediction. By Eq. (3), iCaRL [1] is seen to have $\gamma_{i|k} = 1$ and $\mathcal{D}(p_{i|k}^*, p_{i|k})$ the same as the binary cross-entropy between $p_{i|k}^*$ and $p_{i|k}$, while End-to-End [17] sets $\gamma_{i|k}$ and $\mathcal{D}(p_{i|k}^*, p_{i|k})$ to $p_{i|k}^*$ and $(\log p_{i|k}^* - \log p_{i|k})$, respectively.

Knowledge Distillation as Feature-Graph Preservation. A graph interpretation can help us understand better what we are observing from Eq. (3). We first note that Eq. (3) is to be optimized towards zero for all x_k in X_{all} and is

¹ This amounts to minimizing the cross-entropy $\sum_i \{-p_{i|k}^* \times \log p_{i|k}\}$ between $p_{i|k}^*$ and $p_{i|k}$ over $i \in C_{old}$ since $p_{i|k}^*$ is not affected by the model update.

concerned with the change in relations between the feature vector f_k of certain $x_k \in X_{all}$ and the class embeddings $w_i, \forall i \in C_{old}$. Their relations are seen to be quantified by $p_{i|k}$; specifically, the higher the value of $p_{i|k}$, the more strongly related the f_k and w_i is. This relationship, as depicted in Fig. 2, can be visualized as a bipartite graph, where the feature vectors f_k of all $x_k \in X_{all}$ form a set of vertices that is disjoint from another set of vertices comprising $w_i, i \in C_{old}$. The graph is bipartite because Eq. (3) does not concern the relations among vertices in each of these sets, and each f_k is connected to all the w_i with edge weights specified by $p_{i|k}$. We then view the optimization of the distillation loss over $x_k \in X_{all}$ towards zero as the structure preservation of such a bipartite graph when the learning progresses from one phase to another. This perspective is motivated by the fact that $\mathcal{D}(p_{i|k}^*, p_{i|k})$ in Eq. (3) reflects how the strength of an edge changes as a result of the model update in a learning phase. In particular, edge-adaptive preservation is achieved when $\gamma_{i|k}$ is assigned non-uniformly.

To shed light on the rationale of our design, we now delve deeper into iCaRL [1] and End-to-End [17] in terms of their edge weighting scheme $\gamma_{i|k}$, edge strength specification $p_{i|k}$, and edge discrepancy measure $\mathcal{D}(p_{i|k}^*, p_{i|k})$. First, iCaRL [1] applies equal weighting to all edges on the bipartite graph by having $\gamma_{i|k} = 1, \forall i, k$, suggesting that the strength of every edge should be preserved equally after the model update. In contrast, End-to-End [17] attaches more importance to stronger edges from the viewpoint of the old model by having $\gamma_{i|k} = p_{i|k}^*$. Second, using sigmoid activation in Eq. (1), iCaRL [1] evaluates the edge strength $p_{i|k}$ as a function of the similarity between f_k and w_i alone, whereas the softmax activation (cp. Eq. (2)) in End-to-End [17] renders $p_{i|k}$ dependent on the relative similarities between f_k and all the $w_i, i \in C_{old}$. As such, End-to-End [17] imposes the additional constraint $\sum_i \gamma_{i|k} = \sum_i p_{i|k}^* = 1$ on edge weighting, implying a dependent weighting scheme across edges connecting f_k and $w_i, i \in C_{old}$. Third, the discrepancy measure $\mathcal{D}(p_{i|k}^*, p_{i|k})$ involves $p_{i|k}^*$ and $p_{i|k}$, which in turn depends on the activations $a_{i|k}^*$ and $a_{i|k}$, respectively. These activations reflect explicitly the similarity, or the geometrical relation, between f_k and w_i in the feature space. When viewed as a function of $a_{i|k}$, the $\mathcal{D}(p_{i|k}^*, p_{i|k})$ in both iCaRL [1] and End-to-End [17] is non-symmetrical with respect to $a_{i|k}^*$; that is, the same deviation of $a_{i|k}$ from $a_{i|k}^*$ yet in different signs causes $\mathcal{D}(p_{i|k}^*, p_{i|k})$ to vary differently. This implies that their distillation losses may favor some non-symmetrical geometry variations in the feature space.

Weighted-Euclidean Regularization. In designing the distillation loss, we deviate from the divergence-based criterion to adopt a feature-graph preservation approach. This is motivated by our use of the nearest-mean-of-exemplars classifier, which stresses the geometric structure of the feature space. From the above observations, we conjecture that good regularization for knowledge distillation in the feature space should possess three desirable properties, including (1) prioritized edge preservation, (2) independent edge strength specification, and (3) $a_{i|k}^*$ -symmetric discrepancy measure. The first allows more flexibility in adapting feature learning to new classes while preserving old knowledge, by prioritizing stronger edges. The second dispenses with the constraint

that $\sum_i \gamma_{i|k} = \sum_i p_{i|k}^* = 1$, since the sum-to-one constraint is essential to a cross-entropy interpretation in End-to-End [17], but is not necessary when $p_{i|k}^*$ is regarded as an edge weighting $\gamma_{i|k}$ from our feature-graph viewpoint. The third is to penalize equally the positive and negative deviation from $a_{i|k}^*$ for a geometrically-symmetric preservation of the graph in the feature space. To this end, our *weighted-Euclidean regularization* chooses $a_{i|k}$, $p_{i|k}$, $\gamma_{i|k}$ and $D(p_{i|k}^*, p_{i|k})$ to be

$$\begin{aligned} a_{i|k} &= -\frac{\|\bar{w}_i - \bar{f}_k\|_2^2}{2}, & p_{i|k} &= \frac{1}{Z} \exp\{a_{i|k}\}, & \gamma_{i|k} &= p_{i|k}^* \\ D(p_{i|k}^*, p_{i|k}) &= \left(\log p_{i|k}^* - \log p_{i|k}\right)^2 = \left(a_{i|k}^* - a_{i|k}\right)^2 \end{aligned} \quad (4)$$

where \bar{w}_i, \bar{f}_k are the class embedding and feature vector with rectified cosine normalization (in Sec. 3.2), Z is a normalization constant, and $a_{i|k}^*, p_{i|k}^*$ are the activation and edge strength evaluated with the old model. In particular, we choose the activation $a_{i|k} = -\|\bar{w}_i - \bar{f}_k\|_2^2/2$ to match the criterion of the nearest-mean-of-exemplars classification [1] at test time. It then follows from Eq. (3) that our distillation loss has a form of

$$\mathcal{L}_{dist-wE}(x_k) = \sum_{i \in C_{old}} \left\{ \exp\left(-\frac{\|\bar{w}_i^* - \bar{f}_k^*\|_2^2}{2}\right) \times \left(\|\bar{w}_i^* - \bar{f}_k^*\|_2^2 - \|\bar{w}_i - \bar{f}_k\|_2^2\right)^2 \right\}, \quad (5)$$

where constant terms have been omitted for brevity. Extensive experiments in Sec. 4.3 validate the deceptively simple design of our distillation loss and confirms its superior performance to the seemingly more principled divergence-based methods (e.g. End-to-End [17]).

3.2 Learning New Classes

Learning new classes in the context of incremental learning is often faced with data imbalance. Each new class usually has more training data than an old class since only a limited amount of old data can be stored (as exemplars in the memory). To address this issue, we adopt a two-pronged approach: (1) minimizing binary cross-entropy as the classification loss and (2) applying *rectified cosine normalization* to increase class separation in the feature space.

Classification Loss. In the literature [17,18,1], there are two main choices of the classification loss: cross-entropy and binary cross-entropy. The former is mostly used for multi-class classification while the latter is popular for multi-label classification. In the context of class-incremental learning (a kind of multi-class classification), the study in [18] shows that the model trained with cross-entropy tends to bias in favor of new classes by increasing their class embeddings w_i and bias terms b_i , to take advantage of the imbalanced data distribution for lower classification errors (i.e. the model would quickly adapt to new classes without being able to maintain well the knowledge of old classes). This is because cross-entropy adopts a softmax activation function, which determines the class prediction based on the *relative* strength of the class activations, thereby creating

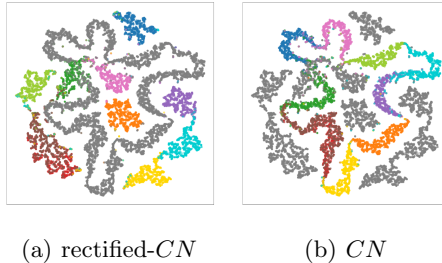


Fig. 3. A toy example based on MNIST for showcasing the benefit of our rectified cosine normalization (denoted as *rectified-CN*) with respect to the typical cosine normalization (denoted as *CN*). The colored dots in **(a)** (respectively **(b)**) are the t-SNE visualization of the feature representations for images of 10 digit classes, which are obtained from the classifier trained with adopting our *rectified-CN* (respectively *CN*); while the gray dots are the features from the classifier of using *CN* (respectively *rectified-CN*).

the aforementioned short cut. With the aim of achieving more evenly distributed classification accuracy, we choose the stricter binary cross-entropy to leverage its property that the activation of each class is weighted *absolutely* rather than *relatively* in determining the class prediction. In symbols, it is evaluated as

$$\mathcal{L}_{bce}(x_k) = - \sum_{i \in C_{all}} \left[\delta_{y_i=y_k} \log(\sigma(a_{i|k})) + \delta_{y_i \neq y_k} \log(1 - \sigma(a_{i|k})) \right], \quad (6)$$

where y_k is the label of the data sample x_k , $C_{all} = C_{old} \cup C_{new}$, and δ is an indicator function.

Rectified Cosine Normalization. In [18], they additionally adopt cosine normalization to address data imbalance by normalizing the class embedding w_i and the feature vector f_k in evaluating the activation $a_{i|k} = \bar{w}_i^T \bar{f}_k$, where \bar{w}_i, \bar{f}_k are the normalized class embedding and feature vector. Here we further develop a *rectified cosine normalization* technique, which is empirically shown to encourage greater separation between incrementally-learned classes in the feature space than the typical cosine normalization. Our rectified cosine normalization is implemented by (1) augmenting, during training with Eq. (6), every class embedding w_i with a learnable bias b_i and every feature vector f_k with a constant 1 for separate normalization; and by (2) involving, at test time, only the feature vector f_k (without augmentation) in the nearest-mean-of-exemplars classification [1]. Step 1 gives rise to an *augmented feature space* of $W_i = (w_i, b_i)$ and $F_k = (f_k, 1)$, with the activation² evaluated as $a_{i|k} = \bar{W}_i^T \bar{F}_k$ where the bar indicates normalization.

Here we use MNIST dataset [26], which is composed of images of handwritten digits, to showcase the benefit of our rectified cosine normalization. We construct

² This activation is to be distinguished from the one $a_{i|k} = -\|\bar{w}_i - \bar{f}_k\|_2^2/2$ for weighted-Euclidean regularization.

the classifier by a shallow network (including 2 convolutional layers followed by 3 fully connected layers with an output size of 3, which implies that the feature vector f is three-dimensional), and train it with whole 10 digit classes. Fig. 3 shows the t-SNE visualization of the feature representations f learned from both the model variants of using our rectified cosine normalization and the typical cosine normalization. It is clear to see that our proposed normalization scheme better encourages the separation between classes. In Sec. 4.3, we present an ablation study to demonstrate the effectiveness of our rectified cosine normalization in improving incremental accuracy.

3.3 Objective Function

We combine the binary cross-entropy classification loss, rectified cosine normalization, and the weighted-Euclidean regularization to build our overall objective:

$$\mathcal{L}(\mathbf{x}) = \frac{1}{N} \sum_{k=1}^N (L_{bce}(x_k) + \lambda \mathcal{L}_{dist_wE}(x_k)), \quad (7)$$

where $a_{i|k} = \bar{W}_i^T \bar{F}_k$ in $\mathcal{L}_{bce}(x_k)$, $\{x_1, x_2, \dots, x_N\}$ is a mini-batch drawn from X_{all} , and λ is the hyper-parameter used to balance \mathcal{L}_{bce} against \mathcal{L}_{dist_wE} . Inspired by [27], we choose λ as follows in recognizing that: in the early stages of incremental learning, when the relative size of old classes to new classes is small, the weighted-Euclidean regularization should be de-emphasized as the feature extractor may not be well learned yet. Note that the ratio $|C_{old}|/|C_{all}|$ is updated constantly along the incremental learning process.

$$\lambda = \lambda_{base} \sqrt{\frac{|C_{old}|}{|C_{all}|}}, \quad (8)$$

where $\lambda_{base} = 0.1$ is set empirically.

4 Experimental Results

Datasets and Baselines. We follow the same experiment protocol as prior works. Two datasets, CIFAR-100 [28] and ImageNet [29] with 100 classes selected randomly, are used for experiments. The data samples of these 100 classes comprise a data stream that arrives in a class-incremental manner in 10 phases, each adding 10 new classes to those already received in previous phases. Particularly, we adopt two training scenarios. One is to train the model from scratch, i.e. the model starts with random initialization and learns 100 classes incrementally. The other follows [18] to train the model with the first 50 classes in order to have a reasonably good feature extractor to begin with. Then the remaining 50 classes are split evenly and learned in 5 sequential phases. The baselines include iCaRL [1], End-to-End [17], BIC [20], Hou19 [18], and Mnemonics [21].

Metrics. The evaluation basically adopts the same common metrics used in prior works, including incremental accuracy and average incremental accuracy [17,18,1]. Additionally, we introduce phase accuracy. *Incremental accuracy* [17], also known simply as accuracy in [1], is the accuracy for classifying all the seen classes at the end of each training phase. It is the most commonly used metric but has the limitation of showing only one single accuracy value over the seen classes as a whole without giving any detail of how the model performs on separate groups of classes (learned incrementally in each phase). *Average incremental accuracy* simply takes the average of the incremental accuracy values obtained from the first training phase up to the current phase. The results are given in parentheses for the last phase in Fig. 4. *Phase accuracy* is evaluated at the end of the entire incremental training to present the average classification accuracy on separate groups of classes, where each group includes several classes that are added into the training at the same phase. It provides a breakdown look at whether the model would favor some groups of classes over the others as a consequence of catastrophic forgetting.

Implementation Details. For a fair comparison, we follow the baselines to use a 32-layer (respectively, 18-layer) ResNet [30] as the feature extractor for CIFAR-100 (respectively, ImageNet), and a fully connected layer as the linear classifier. In particular, we remove the last ReLU layer for rectified cosine normalization. The memory size is fixed at 1000 or 2000, and the exemplars of old classes are chosen by the herd selection [31] proposed in iCaRL [1]. At test time, we apply the same *nearest-mean-of-exemplars* classification strategy as iCaRL.

4.1 Incremental Accuracy Comparison

Fig. 4 presents incremental accuracy against the number of classes learned over the course of incremental learning, while Table 1 summarizes the results at the end of the entire training. Notice that we use five random orderings of classes to build up class-incremental phases, where the metrics are evaluated and averaged over these five orderings. Please also note that the figures showing the experimental results are better viewed in color and zoomed in for details.

We see that on CIFAR-100, when learning from scratch, our model outperforms all the baselines in almost every training phase, showing a significant boost in incremental accuracy especially with memory size 1000. In comparison, the recent state-of-the-arts from Hou19 [18] and Mnemonics [21] perform significantly worse as they rely heavily on a well-learned feature extractor (which is hard to obtain while being trained from scratch). When starting from 50 classes, ours performs the best with memory size 1000, and comparably to both Hou19 [18] and Mnemonics [21] with memory size 2000. Furthermore, our model shows higher average incremental accuracy than the baselines in nearly every setting, as shown in parentheses in Fig. 4. On ImageNet, our model performs comparably to End-to-End [17] and BIC [20] when learning from scratch with $M = 2K$ and is superior to the competing methods by a large margin when learning from 50 classes. To sum up, our method shows consistently higher incremental accuracy than the other baselines on different datasets with varied characteristics.

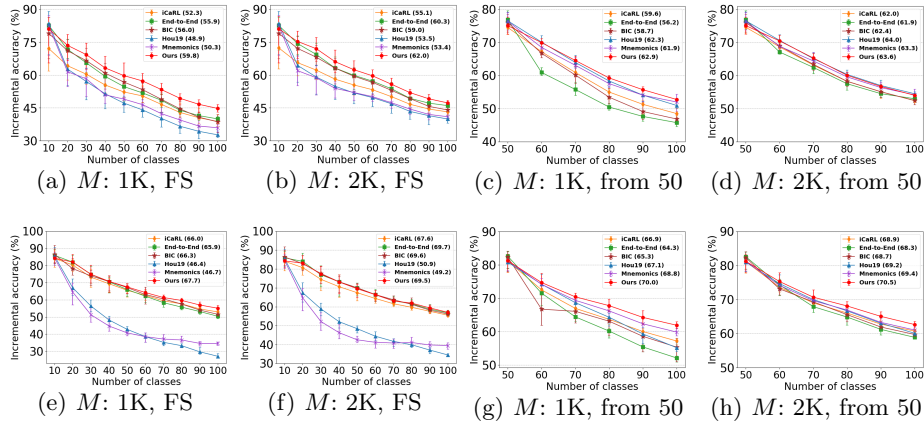


Fig. 4. Incremental accuracy on CIFAR-100 (top row) and ImageNet (bottom row), for memory sizes $M = 1K, 2K$ and training scenarios “FS” and “from 50”. Average incremental accuracy of each method is shown in parentheses.

Table 1. Comparison of incremental accuracy at the end of training

Training scenario	from scratch				from 50 classes			
	CIFAR		ImageNet		CIFAR		ImageNet	
Memory size	1K	2K	1K	2K	1K	2K	1K	2K
iCaRL [1]	38.8	43.1	52.9	55.5	48.5	52.4	57.2	60.7
End-to-End [17]	40.0	46.0	50.4	56.7	45.8	53.0	52.1	58.9
BIC [20]	38.6	44.0	51.5	56.1	46.9	52.4	55.3	59.7
Hou19 [18]	32.7	39.9	27.1	34.4	51.0	54.5	55.2	60.1
Mnemonics [21]	35.9	40.9	34.6	39.4	52.0	54.2	59.8	61.0
Ours	44.7	47.2	55.1	57.0	52.8	54.0	61.9	62.6

4.2 Phase Accuracy Comparison

Fig. 5 presents the phase accuracy for different methods to compare their effectiveness in preserving knowledge of old classes. Generally, balanced phase accuracy is desirable. It is important to point out that there is a fundamental trade-off between incremental accuracy and phase accuracy. For a fair comparison, we particularly choose the training scenario “from 50”, where the baselines perform more closely to our method in terms of incremental accuracy evaluated at the end of the entire training. Shown in parentheses is the mean absolute deviation from the average of each method’s phase accuracy. The smaller the deviation, the more balanced the classification accuracy is in different training phases. Our scheme is shown to achieve the minimum mean absolute deviation in phase accuracy on CIFAR-100 and ImageNet. Remarkably, among all the baselines, Hou19 [18] and Mnemonics [21] have the closest incremental accuracy to ours, yet with considerable variations in phase accuracy.

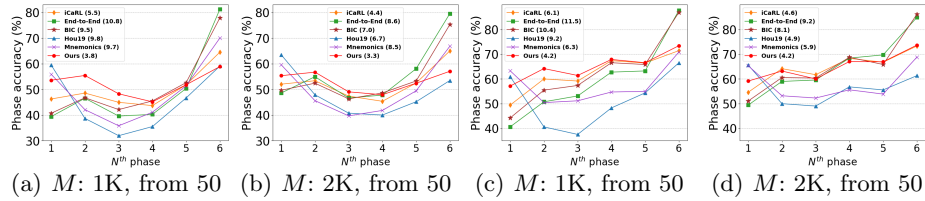


Fig. 5. Phase accuracy comparison: (a)(b) are results on CIFAR-100 and (c)(d) on ImageNet. M is the memory size, and the model is pre-trained with 50 classes. The mean absolute deviation in phase accuracy is shown in parentheses.

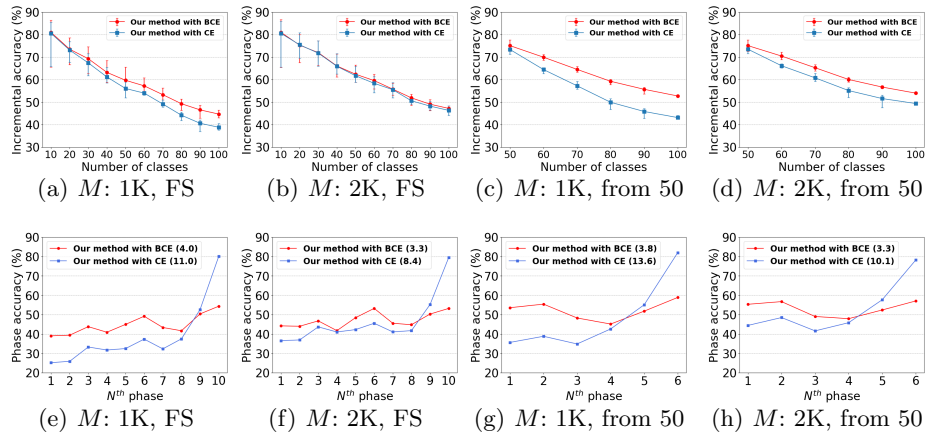


Fig. 6. Comparison of binary cross-entropy and cross-entropy as the classification loss on CIFAR-100 for memory sizes $M = 1K, 2K$ and training scenarios “FS” and “from 50”. The results are evaluated in terms of incremental accuracy. The bars are evaluated by five random orderings of classes.

4.3 Ablation Studies

Binary Cross-entropy (BCE) versus Cross-entropy (CE) as Classification Loss. We observe the performance differences by replacing BCE in our method with CE while keeping the other aspects (including rectified cosine normalization and weighted-Euclidean regularization) untouched. Fig. 6 shows that BCE achieves consistently higher incremental accuracy than CE under all the settings. Moreover, BCE presents more evenly balanced phase accuracy whereas CE displays a tendency of biasing towards those new classes in the last phase.

How effective is rectified cosine normalization in improving incremental accuracy? Fig. 7 presents results comparing rectified cosine normalization against cosine normalization in terms of incremental accuracy. Recall that rectified cosine normalization is proposed as a means to increase class separation and thus classification accuracy. To see the performance differences, we train a

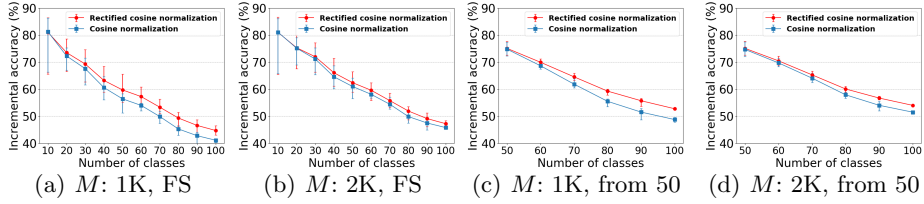


Fig. 7. Comparison between rectified cosine normalization and cosine normalization on CIFAR-100 in terms of incremental accuracy, for memory sizes $M = 1K, 2K$ and training scenarios “FS” and “from 50”.

variant of our method by replacing solely rectified cosine normalization with cosine normalization (while keeping the other aspects the same). From Fig. 7, our rectified cosine normalization starts with similar incremental accuracy to cosine normalization in early phases and gradually develops higher accuracy along the incremental learning process, which validates our analysis in Sec. 3.2.

Efficacy of weighted-Euclidean regularization in balancing new class learning and old knowledge preservation. Fig. 8 compares our weighted-Euclidean regularization \mathcal{L}_{dist_wE} with various distillation losses (see Sec. 3.1 for \mathcal{L}_{dist_bce} and \mathcal{L}_{dist_KL} , and Hou19 [18] for L_{dis}^G) by examining their trade-offs between incremental and phase accuracy. These variants are tested on the same base model of using BCE and rectified cosine normalization. We see that weighted-Euclidean regularization achieves higher incremental accuracy among the competing methods, and shows lowest or the second-lowest mean absolute deviation (MAD) in phase accuracy in most cases (cf. training all classes together in batch mode leads to 2.5 MAD in phase accuracy). These suggest that our weighted-Euclidean regularization enables the model to strike a better balance between new class learning and old knowledge preservation.

Uniform vs. Non-uniform $\gamma_{i|k}$ Assignment. This experiment investigates the sole benefits of prioritized edge preservation (i.e. non-uniform assignment of $\gamma_{i|k}$) in our weighted-Euclidean regularization. We construct a variant of our method by setting $\gamma_{i|k}$ in Eq. (4) to 1 and observe the performance differences. As shown in Fig. 9, the non-uniform assignment of $\gamma_{i|k}$ in our scheme leads to higher incremental accuracy and lower mean absolute deviations in phase accuracy. In addition, the smaller variations in incremental accuracy (cp. the bars in the top row of Fig. 9) suggest that our method with prioritized edge preservation is less sensitive to the learning order of object classes.

5 Conclusion

We introduce a feature-graph preservation perspective to study the essence of various knowledge distillation techniques in how they preserve the structure of the feature graph along the phases of class-incremental learning. Developed based

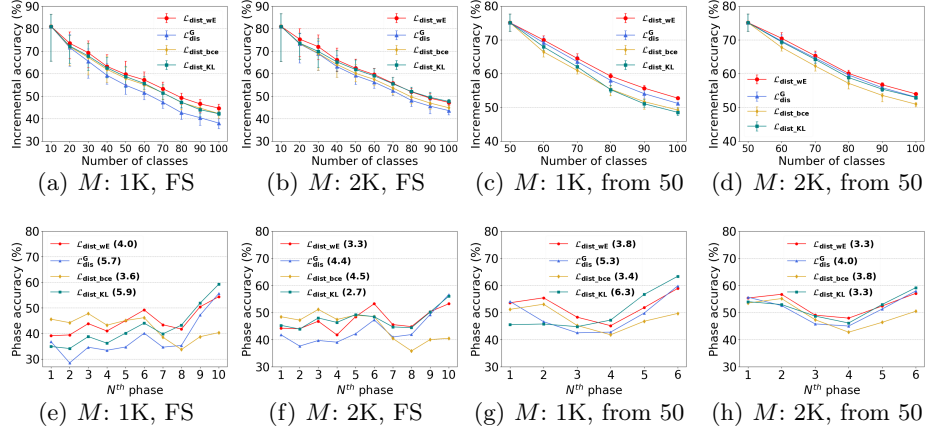


Fig. 8. Comparison of various distillation losses on CIFAR-100, for memory sizes $M = 1K, 2K$ and training scenarios “FS” and “from 50”, with the mean absolute deviation in phase accuracy shown in parentheses.

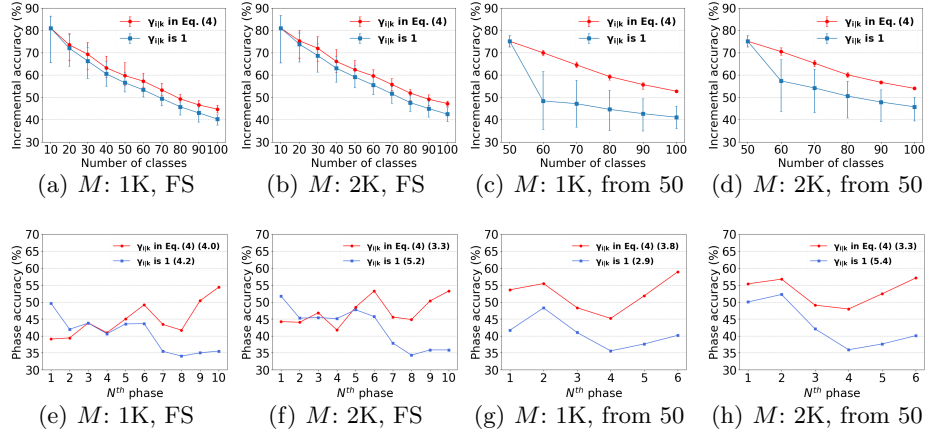


Fig. 9. Comparison of uniform versus non-uniform $\gamma_{i|k}$ assignment on CIFAR-100 for memory sizes $M = 1K, 2K$ and training scenario “FS” and “from 50,” with the mean absolute deviation in phase accuracy (the bottom row) shown in parentheses. The bars in the top row are evaluated by five random orderings of classes.

on three desirable properties, our weighted-Euclidean regularization is shown to benefit the balance between new class learning and old knowledge preservation. Furthermore, our rectified cosine normalization effectively improves incremental accuracy due to greater class separation. With all the techniques combined, our model consistently outperforms the state-of-the-art baselines on CIFAR-100 and ImageNet datasets.

References

1. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017)
2. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: European Conference on Computer Vision (ECCV). (2018)
3. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the National Academy of Sciences (PNAS) (2017)
4. Li, Z., Hoiem, D.: Learning without forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2017)
5. Hou, S., Pan, X., Change Loy, C., Wang, Z., Lin, D.: Lifelong learning via progressive distillation and retrospection. In: European Conference on Computer Vision (ECCV). (2018)
6. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017)
7. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. In: Advances in Neural Information Processing Systems (NIPS). (2017)
8. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient lifelong learning with a-gem. In: International Conference on Learning Representations (ICLR). (2019)
9. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: Psychology of Learning and Motivation. (1989)
10. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: International Conference on Machine Learning (ICML). (2017)
11. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: European Conference on Computer Vision (ECCV). (2018)
12. Aljundi, R., Rohrbach, M., Tuytelaars, T.: Selfless sequential learning. In: International Conference on Learning Representations (ICLR). (2019)
13. Mallya, A., Davis, D., Lazebnik, S.: Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In: European Conference on Computer Vision (ECCV). (2018)
14. Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018)
15. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NIPS Workshop on Deep Learning and Representation Learning. (2014)
16. Belouadah, E., Popescu, A.: Deesil: Deep-shallow incremental learning. In: European Conference on Computer Vision (ECCV). (2018)
17. Castro, F.M., Marín-Jiménez, M.J., Guil, N., Schmid, C., Alahari, K.: End-to-end incremental learning. In: European Conference on Computer Vision (ECCV). (2018)
18. Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019)

19. Jung, H., Ju, J., Jung, M., Kim, J.: Less-forgetful learning for domain expansion in deep neural networks. In: AAAI Conference on Artificial Intelligence (AAAI). (2018)
20. Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., Fu, Y.: Large scale incremental learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019)
21. Liu, Y., Su, Y., Liu, A.A., Schiele, B., Sun, Q.: Mnemonics training: Multi-class incremental learning without forgetting. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020)
22. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Advances in Neural Information Processing Systems (NIPS). (2014)
23. Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. In: Advances in Neural Information Processing Systems (NIPS). (2017)
24. He, C., Wang, R., Shan, S., Chen, X.: Exemplar-supported generative reproduction for class incremental learning. In: British Machine Vision Conference (BMVC). (2018)
25. Ostapenko, O., Puscas, M., Klein, T., Jahnichen, P., Nabi, M.: Learning to remember: A synaptic plasticity driven framework for continual learning. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019)
26. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* (1998)
27. Taitelbaum, H., Ben-Reuven, E., Goldberger, J.: Adding new classes without access to the original training data with applications to language identification. In: Annual Conference of the International Speech Communication Association (INTERSPEECH). (2018)
28. Krizhevsky, A.: Learning multiple layers of features from tiny images. Technical report (2009)
29. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2009)
30. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016)
31. Welling, M.: Herding dynamical weights to learn. In: International Conference on Machine Learning (ICML). (2009)