

Class-Incremental Learning by Knowledge Distillation with Adaptive Feature Consolidation

Minsoo Kang[†]Jaeyoo Park[†]Bohyung Han^{†,§}ECE[†], ASRI[†], & IPAI^{†,§}
Seoul National University

{kminsoo, bellos1203, bhhan}@snu.ac.kr

Abstract

We present a novel class incremental learning approach based on deep neural networks, which continually learns new tasks with limited memory for storing examples in the previous tasks. Our algorithm is based on knowledge distillation and provides a principled way to maintain the representations of old models while adjusting to new tasks effectively. The proposed method estimates the relationship between the representation changes and the resulting loss increases incurred by model updates. It minimizes the upper bound of the loss increases using the representations, which exploits the estimated importance of each feature map within a backbone model. Based on the importance, the model restricts updates of important features for robustness while allowing changes in less critical features for flexibility. This optimization strategy effectively alleviates the notorious catastrophic forgetting problem despite the limited accessibility of data in the previous tasks. The experimental results show significant accuracy improvement of the proposed algorithm over the existing methods on the standard datasets. Code is available.¹

1. Introduction

Deep neural networks have achieved outstanding results in various applications including computer vision [13, 29, 31, 32, 38], natural language processing [45, 47], speech recognition [5, 10], robotics [25], bioinformatics [6], and many others. Despite their impressive performance on offline learning problems, it is still challenging to train models for a sequence of tasks in an online manner, where only a limited number of examples in the previous tasks are available due to memory constraints or privacy issues.

Although fine-tuning is a good strategy to learn a new task given an old model, it is not effective for streaming tasks due to the catastrophic forgetting problem [30]; the

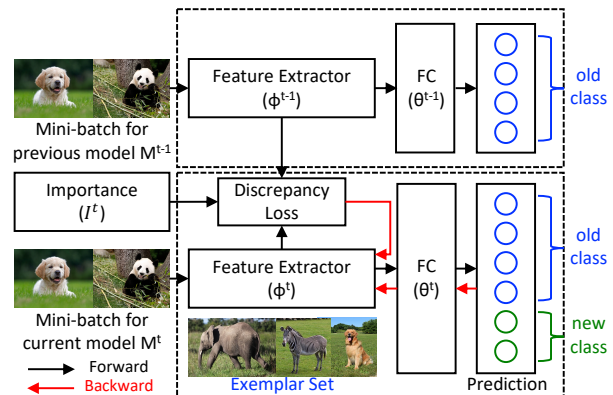


Figure 1. Illustration of the proposed method. The current model optimizes the task loss over a mini-batch sampled from the data of the current task and exemplar sets for the previous tasks while it aims to differently minimize the discrepancy in each feature map over the previous model based on the corresponding importance.

model performs well on the current task while it often fails to generalize on the previous ones. Therefore, incremental learning, the framework capacitating online learning without forgetting, has been studied actively. To mitigate the catastrophic forgetting problem, several branches of algorithms have been discussed. Architectural approaches employ network expansion schemes [27, 41, 49] instead of using static models. Rehearsal methods present the strategies to summarize the tasks in the past by storing exemplar sets [4, 37] or generate samples by estimating the data distribution in the previous tasks [43]. On the other hand, parameter regularization methods [3, 22, 51] prevent important weights from deviating the previously learned models. Knowledge distillation approaches [21, 26] focus on minimizing the divergences from the representations of the previous models while effectively adapting to new tasks.

We present a knowledge distillation approach for class incremental learning, where the model incrementally learns new classes via knowledge distillation, with limited acces-

¹<https://github.com/kminsoo/AFC>

sibility to the data from the classes in the old tasks. Existing approaches based on knowledge distillation simply minimize the distances between the representations of the old and new models without considering which feature maps are important to maintain the previously acquired knowledge. Although [35] addresses this issue, it is limited to the heuristic assignments of the importance weights. On the contrary, we estimate how the representation change given by a model update affects the loss, and show that the minimization of the loss increases can be achieved by a proper weighting of feature maps for knowledge distillation. In a nutshell, the proposed approach maintains important features for robustness, while making less critical features flexible for adaptivity. Note that our optimization strategy aims to minimize the upper bound of the expected loss increases over the previous tasks, which eventually reduces the catastrophic forgetting problem. Figure 1 illustrates the main idea of the proposed approach. The main contributions and characteristics of our algorithm are summarized below:

- We propose a simple but effective knowledge distillation algorithm for class incremental learning based on knowledge distillation with feature map weighting.
- We theoretically show that the proposed approach minimizes the upper bound of the loss increases over the previous tasks, which is derived by recognizing the relationship between the distribution shift in a feature map and the change in the loss.
- Experimental results demonstrate that the proposed technique outperforms the existing methods by large margins in various scenarios.

The rest of the paper is organized as follows. Section 2 discusses related works to class incremental learning. The details of our method is described in Section 3, and the experimental results are presented in Section 4. Finally, we conclude this paper in Section 5.

2. Related Work

Existing class incremental learning approaches are categorized into one of the following four groups: architectural, rehearsal, parameter regularization, and knowledge distillation methods. This section reviews the previous methods in each category.

2.1. Architectural Methods

Most architectural methods adjust network capacity dynamically to handle a sequence of incoming tasks. Progressive networks [41] augment a network component whenever learning a new task. Continual neural Dirichlet process mixture [24] estimates the joint likelihood of a new sample and its label by which it decides whether an additional

network component is required to handle the new sample. There exist a few algorithms related to network augmentation, which increase the number of active channels adaptively with a sparse regularization to maintain the proper model size in a data-driven way [1, 49]. Although these approaches are effective to deal with a long sequence of tasks, they require creating and storing additional network components and performing multiple forward pass computations for inference, which incur extra computational costs.

2.2. Rehearsal and Pseudo-Rehearsal Methods

Rehearsal methods save a small number of examples that represent the whole dataset effectively and utilize the stored data to learn a new task; the techniques in this category often focus on sample selection for rehearsal Incremental Classifier and Representation Learning (iCaRL) [37] selects the examples in a greedy manner with a cardinality constraint by approximating the mean of training data. Also, [4] identifies an exemplar set for the best approximation of the feasible region in the model parameter space induced by the whole data. On the other hand, pseudo-rehearsal techniques estimate the data distribution of the previous tasks and train new models using generated samples together with the real data of the new task to mitigate the class imbalance problem. For example, [34] synthesizes the fake samples of the old tasks to balance the number of examples between old and new classes by employing class-conditional image synthesis networks [33].

2.3. Parameter Regularization Approaches

Parameter regularization approaches encourage or penalize the updates of individual parameters depending on their importance. To be more specific, Elastic Weight Consolidation (EWC) [22] estimates the rigidity or flexibility of each parameter based on the Fisher information matrix and employs its relevance to the previous tasks for model update. Synaptic Intelligence (SI) [51] also provides the weight of each parameter by estimating the path integral along the optimization trajectory for individual parameters. On the other hand, [3] calculates the importance in an unsupervised manner by computing the gradient of the magnitude of the network output. However, as discussed in [16, 46], parameter regularization methods empirically present relatively lower accuracy than the techniques in other categories, and might serve as a poor proxy for keeping the network output [8].

2.4. Knowledge Distillation Approaches

Knowledge distillation is originally proposed to achieve better generalization performance by transferring knowledge from a pretrained model (*i.e.*, teacher) to a target network (*i.e.*, student) through matching their logits [7], output distributions [14], intermediate activations [39], or attention maps [50]. In the class incremental learning scenarios,

the model trained for old tasks is often considered as the teacher network while the student network corresponds to the model to be learned with the additional data in a new task from the initialization point given by the pretrained weights of the teacher.

As the teacher’s knowledge, iCaRL [37] employs the sigmoid output of each class in the previous tasks while [9, 48, 52] use the normalized softmax outputs with a temperature scaling. Unified Classifier Incrementally via Rebalancing (UCIR) [15] tackles the forgetting problem by maximizing the cosine similarity between the features embedded by the teacher and the student. On the other hand, Learning without Memorizing (LwM) [11] minimizes the difference of the gradients given by the highest score classes in the old tasks with respect to intermediate feature maps, which is motivated by [42]. Pooled Outputs Distillation (POD-Net) [12] minimizes the difference of the pooled intermediate features in the height and width directions instead of performing element-wise comparisons; such a relaxed objective using knowledge distillation turns out to be effective for class incremental learning.

Different from the previous works including the closely related method [35], our approach estimates the importance of each feature map for knowledge distillation and provides the theoretical foundation for the adaptive weighting of feature maps. Note that the proposed knowledge distillation algorithm with the adaptive feature map weighting minimizes the loss increases by model updates over the previous tasks.

3. Proposed Algorithm

This section describes the proposed approach based on the knowledge distillation, which alleviates the catastrophic forgetting issue for class incremental learning.

3.1. Problem Formulation

Convolutional neural networks [13, 18] typically contain multiple basic building blocks, each of which is composed of convolutions, batch normalizations [19], and activation functions. A convolutional neural network learner at an incremental stage t , denoted by $M^t(\cdot)$, involves a classifier, $g(\cdot)$, and L building blocks, $\{f_\ell(\cdot); \ell = 1, 2, \dots, L\}$, which is expressed as

$$\begin{aligned} \hat{y} &= M^t(x) = g(h) \\ &= g(f_L \circ \dots \circ f_\ell \circ \dots \circ f_1(x)). \end{aligned} \quad (1)$$

Note that the classifier maps $h \in \mathbb{R}^d$, an embedding vector given by the last building block, to a vector $\hat{y} \in \mathbb{R}^{n^t}$, where n^t denotes the number of the observed classes until the t^{th} task. The ℓ^{th} building block $f_\ell(\cdot)$ also maps an input $Z_{\ell-1}$ to a set of feature maps $Z_\ell = \{Z_{\ell,1}, Z_{\ell,2}, \dots, Z_{\ell,C_\ell}\}$. As mentioned earlier, the network $M^t(\cdot)$ is initialized by the model parameters in the previous stage $M^{t-1}(\cdot)$ except the

additional parameters introduced to handle new classes in the classifier. For notational convenience, we introduce two functions $F_\ell(\cdot)$ and $G_\ell(\cdot)$ splitting the whole model into two subnetworks as follows:

$$M^t(x) = G_\ell(Z_\ell) = G_\ell(F_\ell(x)). \quad (2)$$

At the time of training M^t , it optimizes the following objective function to mitigate catastrophic forgetting issue while learning new classes:

$$\min_{M^t} \mathbb{E}_{(x,y) \sim \mathcal{P}_{\text{data}}^{t-1}} [\mathcal{L}(M^t(x), y)], \quad (3)$$

where $\mathcal{P}_{\text{data}}^{t-1}$ is the data distribution of the past tasks until the $t-1^{\text{st}}$ incremental learning stage, and $\mathcal{L}(\cdot, \cdot)$ is a task-specific loss, *e.g.*, classification loss in this work. However, it is well known that simply optimizing the objective while learning for new classes leads to catastrophic forgetting. Although the discrepancy between two adjacent continual learning models, $M^{t-1}(\cdot)$ and $M^t(\cdot)$, may be small, the amount of forgetting gets more significant as the errors for the old tasks increase with respect to new models.

To analyze the effect of the distribution change from Z_ℓ to Z'_ℓ obtained respectively by $M^{t-1}(\cdot)$ and $M^t(\cdot)$ in the ℓ^{th} layer, we compute the loss given by Z'_ℓ via the first order Taylor approximation as follows:

$$\begin{aligned} \mathcal{L}(G_\ell(Z'_\ell), y) &\approx \\ \mathcal{L}(G_\ell(Z_\ell), y) &+ \sum_{c=1}^{C_\ell} \left\langle \nabla_{Z_{\ell,c}} \mathcal{L}(G_\ell(Z_\ell), y), Z'_{\ell,c} - Z_{\ell,c} \right\rangle_F, \end{aligned} \quad (4)$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product, *i.e.*, $\langle A, B \rangle_F = \text{tr}(A^\top B)$. Note that $\nabla_{Z_{\ell,c}} \mathcal{L}(G_\ell(Z_\ell), y) \in \mathbb{R}^{H_\ell \times W_\ell}$ is given by the backpropagation in $M^{t-1}(\cdot)$. Using Eq. (4), we define the loss increases due to the distribution shift in the c^{th} channel of the ℓ^{th} layer, $\Delta \mathcal{L}(Z'_{\ell,c})$, as

$$\Delta \mathcal{L}(Z'_{\ell,c}) := \left\langle \nabla_{Z_{\ell,c}} \mathcal{L}(G_\ell(Z_\ell), y), Z'_{\ell,c} - Z_{\ell,c} \right\rangle_F. \quad (5)$$

Instead of solving the optimization problem in Eq. (3), we minimize the expected loss increases of old data incurred by the model updates at the t^{th} incremental learning stage, which is given by

$$\min_{M^t} \sum_{\ell=1}^L \sum_{c=1}^{C_\ell} \mathbb{E}_{(x,y) \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\Delta \mathcal{L}(Z'_{\ell,c}) \right]^2. \quad (6)$$

Minimizing Eq. (6) via the standard stochastic gradient descent method requires the additional backpropagation process in $M^{t-1}(\cdot)$ unless we store the gradient of the feature maps with respect to $M^{t-1}(\cdot)$ for each example, which increases computational cost or memory overhead substantially. We will discuss how to address this issue next.

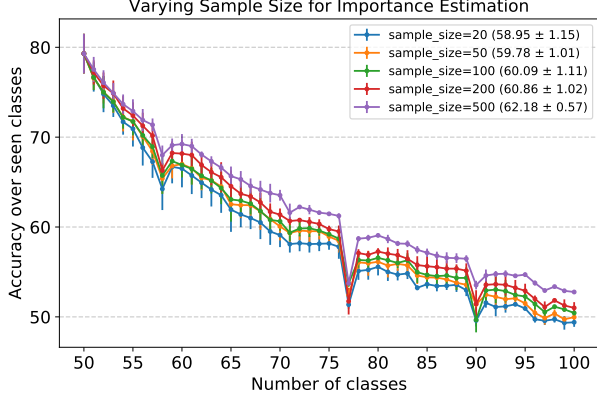


Figure 2. Incremental accuracy on CIFAR100, 50 incremental stages, with ResNet-32 over three different class orders by varying different sample sizes for importance estimation. The legend indicates the mean and standard deviation of the results.

3.2. Practical Objective Function

According to Eq. (6), the distribution shift in each channel affects the loss function with a different magnitude. To avoid storing a large number of feature maps or performing additional backpropagation procedures, we derive an upper bound of the objective function in Eq. (6) as

$$\begin{aligned} & \mathbb{E} \left[\Delta \mathcal{L}(Z'_{\ell,c}) \right] \\ &= \mathbb{E} \left[\langle \nabla_{Z_{\ell,c}} \mathcal{L}(G_{\ell}(Z_{\ell}), y), Z'_{\ell,c} - Z_{\ell,c} \rangle_F \right] \end{aligned} \quad (7)$$

$$\leq \mathbb{E} \left[\|\nabla_{Z_{\ell,c}} \mathcal{L}(G_{\ell}(Z_{\ell}), y)\|_F \cdot \|Z'_{\ell,c} - Z_{\ell,c}\|_F \right] \quad (8)$$

$$\leq \sqrt{\mathbb{E} \left[\|\nabla_{Z_{\ell,c}} \mathcal{L}(G_{\ell}(Z_{\ell}), y)\|_F^2 \right] \cdot \mathbb{E} \left[\|Z'_{\ell,c} - Z_{\ell,c}\|_F^2 \right]}, \quad (9)$$

where the expectations are taken over $\mathcal{P}_{\text{data}}^{t-1}$. Refer to the supplementary file for the detailed derivation.

Then, we define a more practical objective function, the upper bound of Eq. (6), as follows:

$$\min_{M^t} \sum_{\ell=1}^L \sum_{c=1}^{C_{\ell}} I_{\ell,c}^t \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\|Z'_{\ell,c} - Z_{\ell,c}\|_F^2 \right], \quad (10)$$

where the importance is given by

$$I_{\ell,c}^t := \mathbb{E}_{(x,y) \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\|\nabla_{Z_{\ell,c}} \mathcal{L}(G_{\ell}(Z_{\ell}), y)\|_F^2 \right],$$

which is estimated by a Monte Carlo integration at the previous stage using $M^{t-1}(\cdot)$ and employed for training M^t . Note that we need marginal cost for storing $I_{\ell,c}^t$ since it is just a scalar value corresponding to each channel. The importance $I_{\ell,c}^t$ can be interpreted as a weight factor because the larger value of $I_{\ell,c}^t$ incurs the bigger changes in the loss

given by the same perturbation in the feature map. Furthermore, Figure 2 shows that a larger sample sizes estimates more reliable importance and consequently leads to better generalization performance.

3.3. More Robust Objective Function

Although the optimization of the objective function in Eq. (10) is effective to maintain the representations of the examples in the previous tasks, we usually have a limited number of exemplars and face critical challenge in reducing catastrophic forgetting. Therefore, we also adopt the data in the current task together with the ones in the previous tasks, which leads to the following optimization problem:

$$\min_{M^t} \sum_{\ell=1}^L \sum_{c=1}^{C_{\ell}} I_{\ell,c}^t \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^t} \left[\|Z'_{\ell,c} - Z_{\ell,c}\|_F^2 \right], \quad (11)$$

where its only difference from Eq. (10) is the distribution for the computation of the expectation. Note that the use of the examples in the current task for the purpose of preserving the representations in the previous stages is also motivated by other knowledge distillation approaches [11, 12, 15, 26]. This strategy is also supported by the following proposition.

Proposition 1. Let $\mathcal{Y}_{\text{old}}^t$, $\mathcal{Y}_{\text{new}}^t$, and q_{data}^t be a set of previously observed classes (\mathcal{Y}^{t-1}), the set of newly observed classes at t^{th} incremental learning stage, and the distribution of the data in the current task, respectively. Let the distribution of labels in \mathcal{Y}^t be a mixture of two prior distributions corresponding to $\mathcal{Y}_{\text{old}}^t$ and $\mathcal{Y}_{\text{new}}^t$ with weights ϕ_{old}^t and ϕ_{new}^t , respectively. Assuming that the data generating process for each class, $\mathcal{P}_{\text{data}}^t(x|y)$, does not change over the incremental stages, the objective in Eq. (10) with the weight factor ϕ_{old}^t is bounded above by that in Eq. (11).

Proof. By the definition of the label distribution for \mathcal{Y}^t , the class prior distribution $\mathcal{P}_{\text{data}}^t(y)$ is given by

$$\begin{aligned} \mathcal{P}_{\text{data}}^t(y) &= \mathcal{P}_{\text{data}}^t(y|\mathcal{Y}_{\text{old}}^t) \phi_{\text{old}}^t + \mathcal{P}_{\text{data}}^t(y|\mathcal{Y}_{\text{new}}^t) \phi_{\text{new}}^t \\ &= \mathcal{P}_{\text{data}}^{t-1}(y) \phi_{\text{old}}^t + q_{\text{data}}^t(y) \phi_{\text{new}}^t. \end{aligned} \quad (12)$$

Based on this equation, $\mathcal{P}_{\text{data}}^t(x)$ can be expressed as

$$\begin{aligned} \mathcal{P}_{\text{data}}^t(x) &= \sum_{y \in \mathcal{Y}_{\text{old}}^t} \mathcal{P}_{\text{data}}^t(y) \mathcal{P}_{\text{data}}^t(x|y) + \sum_{y \in \mathcal{Y}_{\text{new}}^t} \mathcal{P}_{\text{data}}^t(y) \mathcal{P}_{\text{data}}^t(x|y) \\ &= \phi_{\text{old}}^t \cdot \sum_{y \in \mathcal{Y}^{t-1}} \mathcal{P}_{\text{data}}^{t-1}(y) \mathcal{P}_{\text{data}}^{t-1}(x|y) + \\ &\quad \phi_{\text{new}}^t \cdot \sum_{y \in \mathcal{Y}_{\text{new}}^t} q_{\text{data}}^t(y) q_{\text{data}}^t(x|y) \\ &= \phi_{\text{old}}^t \cdot \mathcal{P}_{\text{data}}^{t-1}(x) + \phi_{\text{new}}^t \cdot q_{\text{data}}^t(x). \end{aligned} \quad (13)$$

Then, by denoting $\Delta Z_{\ell,c} := \|Z'_{\ell,c} - Z_{\ell,c}\|_F^2$, which is always non-negative, we obtain the following inequality:

$$\begin{aligned} & \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\Delta Z_{\ell,c} \right] \\ & \leq \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\Delta Z_{\ell,c} \right] + \frac{\phi_{\text{new}}^t}{\phi_{\text{old}}^t} \mathbb{E}_{x \sim q_{\text{data}}^t} \left[\Delta Z_{\ell,c} \right]. \end{aligned} \quad (14)$$

By multiplying both sides of Eq. (14) by ϕ_{old}^t and using Eq. (13), we have

$$\begin{aligned} & \phi_{\text{old}}^t \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\Delta Z_{\ell,c} \right] \\ & \leq \phi_{\text{old}}^t \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^{t-1}} \left[\Delta Z_{\ell,c} \right] + \phi_{\text{new}}^t \mathbb{E}_{x \sim q_{\text{data}}^t} \left[\Delta Z_{\ell,c} \right] \\ & = \mathbb{E}_{x \sim \mathcal{P}_{\text{data}}^t} \left[\Delta Z_{\ell,c} \right], \end{aligned} \quad (15)$$

which concludes the proof. \square

Proposition 1 implies that the optimization procedure of Eq. (11) minimizes the upper bound of the loss increases over the previous tasks using the examples in the previous and current tasks effectively.

3.4. Final Loss Function

The final goal of $M^t(\cdot)$ is to minimize the combination of the classification loss and the discrepancy loss as

$$\mathcal{L}_{\text{total}}^t = \mathcal{L}_{\text{cls}}^t + \lambda_{\text{disc}} \cdot \lambda^t \mathcal{L}_{\text{disc}}^t, \quad (16)$$

where λ_{disc} is a hyperparameter and an adaptive weight, λ^t , is set to $\sqrt{\frac{n^t}{n^t - n^{t-1}}}$ as proposed in [12, 15]. For the classification loss $\mathcal{L}_{\text{cls}}^t$, we adopt a local similarity classifier (LSC) [12] by adaptively aggregating the cosine similarities obtained from multiple class embedding vectors, which is given by

$$\hat{y}_k = \sum_j s_{k,j} \langle \theta_{k,j}, h \rangle, \quad s_{k,j} = \frac{\exp \langle \theta_{k,j}, h \rangle}{\sum_i \exp \langle \theta_{k,i}, h \rangle}, \quad (17)$$

where $\theta_{k,j}$ is the j^{th} normalized class embedding vector for the k^{th} class and \hat{y}_k is the score for the k^{th} class. Then, the classification loss $\mathcal{L}_{\text{cls}}^t$ is defined with a label $g^{(b)}$ as

$$\mathcal{L}_{\text{cls}}^t = \frac{1}{B} \sum_{b=1}^B \left[-\log \frac{\exp(\eta(\hat{y}_{g^{(b)}} - \delta))}{\sum_{i \neq g^{(b)}} \exp(\eta \hat{y}_i)} \right]_+, \quad (18)$$

where η is a learnable scaling parameter, δ is a constant to encourage a larger inter-class separation, B is the mini-batch size, and $[\cdot]_+$ denotes the ReLU activation function. Also, the discrepancy loss, $\mathcal{L}_{\text{disc}}^t$, is given by

$$\mathcal{L}_{\text{disc}}^t := \frac{1}{B} \sum_{b=1}^B \sum_{\ell=1}^L \sum_{c=1}^{C_\ell} \tilde{I}_{\ell,c}^t \|Z'_{\ell,c}^{(b)} - Z_{\ell,c}^{(b)}\|_F^2, \quad (19)$$

Algorithm 1 Normalized Importance Estimation

- 1: **Input:** current model $M^t(\cdot)$, current dataset D^t , exemplar sets for the previous tasks E^t
 - 2: **Initialization:** $I_{\ell,c}^t \leftarrow 0$ and $D \leftarrow D^t \cup E^t$
 - 3: **for all** $(x, y) \in D$ **do**
 - 4: Compute $\mathcal{L}_{\text{cls}}^t$ by Eq. (18)
 - 5: Perform backward through $M^t(\cdot)$
 - 6: **for** $\ell \in 1, 2, \dots, L$ **do**
 - 7: Update the importance:
 - 8: $I_{\ell,c}^t \leftarrow I_{\ell,c}^t + \|\nabla_{Z_{\ell,c}} \mathcal{L}_{\text{cls}}^t\|_F^2$
 - 9: **end for**
 - 10: **end for**
 - 11: **for** $\ell \in 1, 2, \dots, L$ **do**
 - 12: Normalize the importance $I_{\ell,c}^t$ by Eq. (20)
 - 13: **end for**
 - 14: **Output:** Normalized importance $\{\tilde{I}_{\ell,c}^t\}$ employed for training $M^{t+1}(\cdot)$
-

where $\tilde{I}_{\ell,c}^t$ is introduced to balance the scale of the importances across layers, which is given by

$$\tilde{I}_{\ell,c}^t = \frac{I_{\ell,c}^t}{\frac{1}{C_\ell} \sum_{c=1}^{C_\ell} I_{\ell,c}^t}. \quad (20)$$

Algorithm 1 presents the details of the importance estimation procedure after training $M^t(\cdot)$.

3.5. Management of Exemplar Set

To maintain the exemplar set for each class after training $M^t(\cdot)$, we employ the nearest-mean-of-exemplars selection rule [37] as in other class incremental learning methods. For inference, we employ two methods; one is the nearest-mean-of-exemplars classification rule [37] and the other is based on the classifier probabilities as discussed in [12, 15], which we call NME and CNN, respectively.

3.6. Discussion about Computational Cost

As in most approaches based on knowledge distillation, we perform 2 forward passes and 1 backward pass during training. Different from the previous approaches, at the end of the training phase of each task, we update the importance by performing a pair of forward and backward pass computations using the old exemplars and examples in the current task as presented in Algorithm 1. Therefore, the additional computation of our method is only for the importance update, which happens only once per task; it is negligible compared to the training cost of each task.

4. Experiments

This section compares our algorithm with the previous approaches on the standard datasets, and analyzes the experimental results through ablation studies.

Table 1. Performance comparison between the proposed algorithm, denoted by AFC, and other state-of-the-art methods on CIFAR100. We run experiments using three different class orders and report their averages and standard deviations. Methods with an asterisk * report the results from our reproductions with their official codes. Red and black bold-faced numbers represent the best and second-best performance in each column.

New classes per stage	CIFAR100			
	50 stages 1	25 stages 2	10 stages 5	5 stages 10
iCaRL [37]	44.20 ± 0.98	50.60 ± 1.06	53.78 ± 1.16	58.08 ± 0.59
BiC [48]	47.09 ± 1.48	48.96 ± 1.03	53.21 ± 1.01	56.86 ± 0.46
UCIR (NME) [15]	48.57 ± 0.37	56.82 ± 0.19	60.83 ± 0.70	63.63 ± 0.87
UCIR (CNN) [15]	49.30 ± 0.32	57.57 ± 0.23	61.22 ± 0.69	64.01 ± 0.91
Mnemonics [28]	-	60.96 ± 0.72	62.28 ± 0.43	63.34 ± 0.62
GDumb* [36]	59.76 ± 1.49	59.97 ± 1.51	60.24 ± 1.42	60.70 ± 1.53
PODNet (NME) [12]	61.40 ± 0.68	62.71 ± 1.26	64.03 ± 1.30	64.48 ± 1.32
PODNet (CNN) [12]	57.98 ± 0.46	60.72 ± 1.36	63.19 ± 1.16	64.83 ± 0.98
PODNet (NME)* [12]	56.78 ± 0.41	59.54 ± 0.66	63.27 ± 0.69	65.32 ± 0.65
PODNet (CNN)* [12]	57.86 ± 0.38	60.51 ± 0.62	62.78 ± 0.78	64.62 ± 0.65
GeoDL* [44]	52.28 ± 3.91	60.21 ± 0.46	63.61 ± 0.81	65.34 ± 1.05
SSIL* [2]	53.64 ± 0.77	58.02 ± 0.79	61.52 ± 0.44	63.02 ± 0.59
AFC (NME)	62.58 ± 1.02	64.06 ± 0.73	64.29 ± 0.92	65.82 ± 0.88
AFC (CNN)	62.18 ± 0.57	63.89 ± 0.93	64.98 ± 0.87	66.49 ± 0.81

4.1. Dataset

We employ the standard datasets, CIFAR100 [23] and ImageNet100/1000 [40], for evaluation of class incremental learning algorithms. The CIFAR100 dataset contains 50,000 and 10,000 32×32 color images for training and testing. The ImageNet1000 dataset consists of 1,281,167 images for training and 50,000 images for validation across 1,000 classes. ImageNet100 is a subset of ImageNet1000 with randomly selected 100 classes. For image preprocessing and class order, we follow the protocol in PODNet [12] for fair comparisons. Except the initial task containing half of the total classes for all datasets, the remaining classes are evenly distributed across tasks. After each incremental stage, we store 20 images per class and construct exemplar sets unless stated otherwise.

4.2. Implementation Details

The proposed method is implemented based on the publicly available official code of PODNet and tested using a NVIDIA Tesla V100 GPU, without modifying any experimental settings and network structures for fair comparisons. For all datasets, we adopt the SGD optimizer with a momentum of 0.9, an initial learning rate of 0.1 with a cosine annealing schedule, and a batch size of 128. We train a 32-layer ResNet [13] for 160 epochs with a weight decay of 0.0005 for CIFAR100, and an 18-layer ResNet [13] for 90 epochs with a weight decay of 0.0001 for ImageNet datasets, where all compared methods also use the same backbones. Following the heuristics proposed in [12, 50], we normalize each feature map using its Frobenius norm

to compute the discrepancy loss to further improves performance. We select the same layers with PODNet for knowledge distillation. In case of λ_{disc} , we set it to 4.0 for CIFAR100 and 10.0 for ImageNet.

For comparisons, we report *average incremental accuracy* [37]—the average of the accuracies over seen classes across all incremental stages, where the results of other methods except GeoDL [44], DDE [17], and SSIL [2] are retrieved from [12] if not specified otherwise. Our algorithm is referred to as Adaptive Feature Consolidation (AFC) to present experimental results in this section.

4.3. Results on CIFAR100

We test the proposed method, AFC, in various class incremental learning scenarios. Table 1 clearly shows that the proposed approach achieves the highest accuracy in all the tested numbers of incremental stages. Note that the reproduced results of PODNet with NME inference are worse than the reported ones, even though we use their official code. Compared with the reproduced PODNet and the other techniques, the performance gap becomes larger as the number of stages grows, which implies that the proposed method is more robust to challenging scenarios than the previous methods. AFCs with NME and CNN for inference outperform the reported results of PODNet in [12] and achieve the state-of-the-art performance.

4.4. Results on ImageNet

Table 2 presents the results from all algorithms including AFC on ImageNet100 and ImageNet1000 our algorithm; AFC shows outstanding performance compared to all the

Table 2. Performance comparison between AFC and other state-of-the-art algorithms on ImageNet100 and ImageNet1000.

New classes per stage	ImageNet100				ImageNet1000	
	50 stages 1	25 stages 2	10 stages 5	5 stages 10	10 stages 50	5 stages 100
iCaRL [37]	54.97	54.56	60.90	65.56	46.72	51.36
BiC [48]	46.49	59.65	65.14	68.97	44.31	45.72
UCIR (NME) [15]	55.44	60.81	65.83	69.07	59.92	61.56
UCIR (CNN) [15]	57.25	62.94	70.71	71.04	61.28	64.34
UCIR (CNN) + DDE [17]	-	-	70.20	72.34	65.77	67.51
Mnemonics [28]	-	69.74	71.37	72.58	63.01	64.54
PODNet (CNN) [12]	62.48	68.31	74.33	75.54	64.13	66.95
PODNet (CNN) [12] + DDE [17]	-	-	75.41	76.71	64.71	66.42
GeoDL [44]	-	71.72	73.55	73.87	64.46	65.23
AFC (CNN)	72.08	73.34	75.75	76.87	67.02	68.90

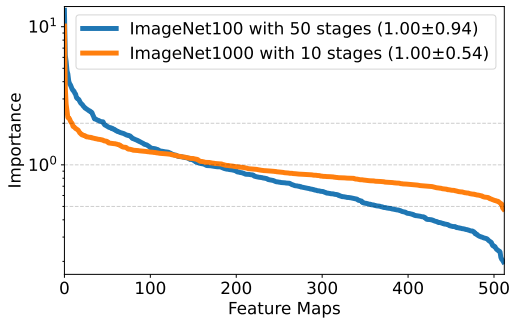


Figure 3. The sorted importance of the feature maps in the last block on ImageNet100 with 50 stages, and ImageNet1000 with 10 stages.

Table 3. Sensitivity analysis regarding the memory budget per class on CIFAR100 with 50 stages. AFC with Eq. (10) optimizes the discrepancy loss over the only exemplar sets without using the data in the current task.

Memory budget per class	5	10	20	50
iCaRL [37]	16.44	28.57	44.20	48.29
BiC [48]	20.84	21.97	47.09	55.01
UCIR (CNN) [15]	22.17	42.70	49.30	57.02
PODNet (CNN) [12]	35.59	48.54	57.98	63.69
AFC (CNN) with (10)	40.37	55.24	61.73	65.10
AFC (CNN)	44.66	55.78	62.18	65.07

compared methods. Moreover, AFC exceeds the previous methods by large margins for 50 stages on ImageNet100, which is more prone to suffer from catastrophic forgetting. These results imply that the proposed approach has great potential for large-scale problems.

4.5. Analysis and Ablation Studies

We analyze the estimated importance and discuss the effects of memory size for exemplar sets, feature distillation

Table 4. Performance comparison of feature distillation approaches on CIFAR100 with 50 stages.

Inference methods	NME	CNN
Fine-tuning	41.56	40.76
Uniform Importance	42.21	40.81
GradCam Preserving [11]	41.89	42.07
Gram Matrix Preserving [20]	41.74	40.80
Channel Preserving [12]	55.91	50.34
Gap Preserving [12]	57.25	53.87
PODNet [12]	61.42	57.64
PODNet* [12]	57.15	57.51
AFC	62.58	62.18

methods, varying initial task sizes, and discrepancy loss.

Importance visualization Figure 3 visualizes the importances of individual feature maps in the last building block, which are sorted after training the final task. It demonstrates that the standard deviation on ImageNet1000 is smaller than that on ImageNet100. This is because more features are needed to effectively learn the knowledge from the examples in ImageNet1000 since the dataset is more challenging. On the other hand, although it is not shown in the figure, we also observe the relatively higher standard deviation in the early stages for both datasets, which is partly because the features in the early stages are less critical to the loss.

Effect of memory budget To assess the effectiveness of the proposed method, we vary the memory budget. Table 3 illustrates that AFC outperforms the state-of-the-art methods by large margins under several different memory budgets. AFC optimized using the exemplar sets together with the current data tends to achieve better generalization performance than AFC with Eq. (10) when the memory budget is small. However, when the budget per class increases as large as 50, AFC with Eq. (10) achieves slightly better accuracy, which is probably because Eq. (10) leads to a tighter

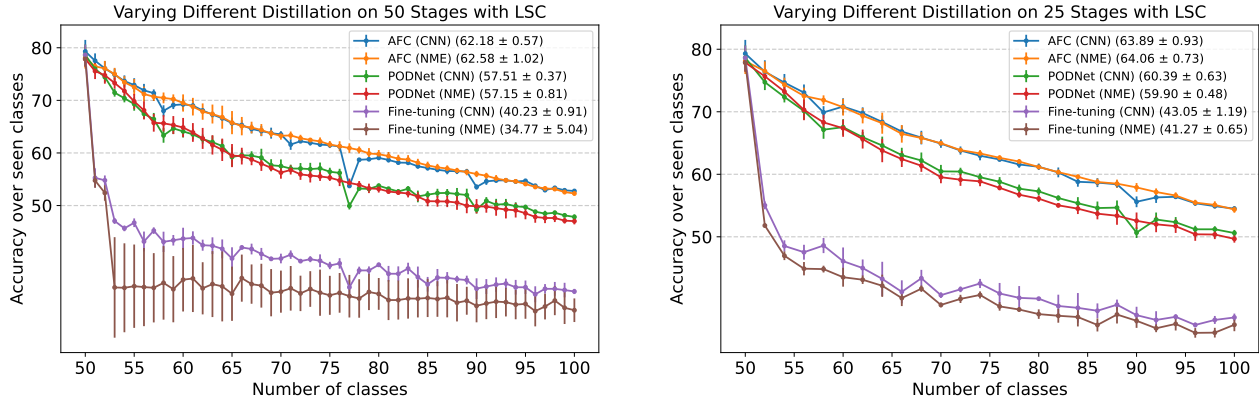


Figure 4. Average accuracies of three different class orders during class incremental learning on CIFAR100 with 50 and 25 stages.

Table 5. Performance comparison between AFC and the state-of-the-art algorithms on CIFAR100 by varying the number of classes in the initial task while each of the remaining tasks only contains a single class.

Initial task size	CIFAR100			
	80 stages	70 stages	60 stages	50 stages
	20	30	40	50
iCaRL [37]	41.28	43.38	44.35	44.20
BiC [48]	40.95	42.27	45.18	47.09
UCIR (NME) [15]	40.81	46.80	46.71	48.57
UCIR (CNN) [15]	41.69	47.85	47.51	49.30
GDumb* [36]	52.04	55.15	57.44	59.76
PODNet (NME) [12]	49.03	55.30	57.89	61.40
PODNet (CNN) [12]	47.68	52.88	55.42	57.98
PODNet (NME)* [12]	45.88	52.18	54.22	56.78
PODNet (CNN)* [12]	48.37	53.28	55.86	57.86
AFC (NME)	51.31	57.05	60.06	62.58
AFC (CNN)	52.90	57.61	60.27	62.18

upper bound and the use of the examples in the current task is not particularly helpful.

Comparison with feature distillation approaches We present an experimental result by applying different feature distillation approaches on CIFAR100 with 50 stages. Table 4 shows the performance of class incremental learning algorithms based on knowledge distillation, where AFC is compared with 7 different methods including “Fine-tuning” and “Uniform Importance” with all the importances set to 1. According to the results, AFC outperforms all the compared algorithms. In addition, Figure 4 also implies that AFC is more robust to catastrophic forgetting than PODNet consistently throughout the incremental stages.

Results by varying initial task sizes We test the proposed method by reducing the number of classes in the initial task while each of the remaining tasks consists of

Table 6. Sensitivity analysis of λ_{disc} on CIFAR100 with 50 stages.

λ_{disc}	2.0	3.0	4.0	5.0	6.0
AFC (NME)	60.16	62.27	62.58	62.37	61.83
AFC (CNN)	59.98	61.88	62.18	62.34	62.08

only one class, which makes the problem more challenging by increasing the total number of the incremental stages. Table 5 presents that AFC outperforms the previous algorithms consistently.

Effect of λ_{disc} for discrepancy loss Table 6 presents the results on CIFAR100 with 50 stages by varying λ_{disc} , which controls the balance between the previous knowledge and the current task. As shown in Table 6, the proposed method is not sensitive to the hyperparameter.

5. Conclusion

We presented a simple but effective knowledge distillation approach via adaptive feature consolidation, which adjusts the weight of each feature map for balancing between adaptivity to new data and robustness to old ones in class incremental learning scenarios. We formulated the importance of a feature in a principled way by deriving the relationship between the discrepancy in the distribution of each feature and the change in the loss. We further relaxed the original objective function and incorporated knowledge distillation to effectively reduce computational cost and memory overhead. Experimental results show that the proposed approach outperforms the existing ones by large margins.

Acknowledgments This work was partly supported by Samsung Electronics Co., Ltd., the IITP grants [2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), 2021-0-02068, Artificial Intelligence Innovation Hub], and the National Research Foundation of Korea (NRF) grant [2022R1A2C3012210] funded by the Korea government (MSIT).

References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Babak Ehteshami Bejnordi. Conditional Channel Gated Networks for Task-Aware Continual Learning. In *CVPR*, 2020. 2
- [2] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. SS-IL: Separated Softmax for Incremental Learning. In *ICCV*, 2021. 6
- [3] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory Aware Synapses: Learning what (not) to forget. In *ECCV*, 2018. 1, 2
- [4] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient Based Sample Selection for Online Continual Learning. In *NeurIPS*, 2019. 1, 2
- [5] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep Speech 2: End-to-End Speech Recognition in English and Mandarin. In *ICML*, 2016. 1
- [6] Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle. Deep Learning for Computational Biology. *Molecular Systems Biology*, 2016. 1
- [7] Jimmy Ba and Rich Caruana. Do Deep Nets really Need to be Deep? In *NIPS*, 2014. 2
- [8] Ari Benjamin, David Rolnick, and Konrad Kording. Measuring and Regularizing Networks in Function Space. In *ICLR*, 2019. 2
- [9] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-End Incremental Learning. In *ECCV*, 2018. 3
- [10] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. Listen, Attend and Spell: A Neural Network for Large Vocabulary Conversational Speech Recognition. In *ICASSP*, 2016. 1
- [11] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without Memorizing. In *CVPR*, 2019. 3, 4, 7
- [12] Arthur Douillard, Matthieu Cord, Charles Ollion, and Thomas Robert. PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning. In *ECCV*, 2020. 3, 4, 5, 6, 7, 8
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 1, 3, 6
- [14] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [15] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a Unified Classifier Incrementally via Rebalancing. In *CVPR*, 2019. 3, 4, 5, 6, 7, 8
- [16] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. *arXiv preprint arXiv:1810.12488*, 2018. 2
- [17] Xinting Hu, Kaihua Tang, Chunyan Miao, Xian-Sheng Hua, and Hanwang Zhang. Distilling Causal Effect of Data in Class-Incremental Learning. In *CVPR*, 2021. 6, 7
- [18] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely Connected Convolutional Networks. In *CVPR*, 2017. 3
- [19] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 2015. 3
- [20] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In *ECCV*, 2016. 7
- [21] Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetting Learning in Deep Neural Networks. *arXiv preprint arXiv:1607.00122*, 2016. 1
- [22] 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*, 2017. 1, 2
- [23] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Learning Multiple Layers of Features from Tiny Images. Technical report, 2009. 6
- [24] Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A Neural Dirichlet Process Mixture Model for Task-Free Continual Learning. In *ICLR*, 2019. 2
- [25] Ian Lenz, Honglak Lee, and Ashutosh Saxena. Deep Learning for Detecting Robotic Grasps. *International Journal of Robotics Research*, 2015. 1
- [26] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *TPAMI*, 2017. 1, 4
- [27] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive Aggregation Networks for Class-Incremental Learning. In *CVPR*, 2021. 1
- [28] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics Training: Multi-Class Incremental Learning without Forgetting. In *CVPR*, 2020. 6, 7
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *CVPR*, 2015. 1
- [30] Michael McCloskey and Neal J Cohen. Catastrophic Interference in Connectionist Networks: The sequential Learning Problem. In *Psychology of Learning and Motivation*. Elsevier, 1989. 1
- [31] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 1
- [32] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *ICCV*, 2015. 1
- [33] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional Image Synthesis with Auxiliary Classifier Gans. In *ICML*, 2017. 2
- [34] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to Remember: A Synaptic Plasticity Driven Framework for Continual Learning. In *CVPR*, 2019. 2

- [35] Jaeyoo Park, Minsoo Kang, and Bohyung Han. Class-incremental learning for action recognition in videos. In *ICCV*, 2021. 2, 3
- [36] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. GDumb: A Simple Approach that Questions Our Progress in Continual Learning. In *ECCV*, 2020. 6, 8
- [37] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental Classifier and Representation Learning. In *CVPR*, 2017. 1, 2, 3, 5, 6, 7, 8
- [38] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. In *CVPR*, 2016. 1
- [39] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. FitNets: Hints For Thin Deep Nets. In *ICLR*, 2015. 2
- [40] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 6
- [41] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive Neural Networks. *arXiv preprint arXiv:1606.04671*, 2016. 1, 2
- [42] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-Cam: Visual Explanations from Deep Networks via Gradient-based Localization. In *ICCV*, 2017. 3
- [43] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual Learning with Deep Generative Replay. In *NIPS*, 2017. 1
- [44] Christian Simon, Piotr Koniusz, and Mehrtash Harandi. On Learning the Geodesic Path for Incremental Learning. In *CVPR*, 2021. 6, 7
- [45] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to Sequence Learning with Neural Networks. In *NIPS*, 2014. 1
- [46] Gido M van de Ven and Andreas S Tolias. Three Scenarios for Continual Learning. *arXiv preprint arXiv:1904.07734*, 2019. 2
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *NIPS*, 2017. 1
- [48] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large Scale Incremental Learning. In *CVPR*, 2019. 3, 6, 7, 8
- [49] Shipeng Yan, Jiangwei Xie, and Xuming He. DER: Dynamically Expandable Representation for Class Incremental Learning. In *CVPR*, 2021. 1, 2
- [50] Sergey Zagoruyko and Nikos Komodakis. Paying More Attention to Attention: Improving the Performance of Convolutional Neural Networks via Attention Transfer. In *ICLR*, 2017. 2, 6
- [51] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning Through Synaptic Intelligence. In *ICML*, 2017. 1, 2
- [52] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining Discrimination and Fairness in Class incremental Learning. In *CVPR*, 2020. 3