

Class-Incremental Learning with Strong Pre-trained Models

Tz-Ying Wu^{1,2} Gurumurthy Swaminathan¹ Zhizhong Li¹
 Avinash Ravichandran¹ Nuno Vasconcelos² Rahul Bhotika¹ Stefano Soatto¹
¹AWS AI Labs ²UC San Diego

{gurumurs,lzhizhon,ravinash,bhotikar,soattos}@amazon.com

{tzw001,nuno}@ucsd.edu

Abstract

Class-incremental learning (CIL) has been widely studied under the setting of starting from a small number of classes (base classes). Instead, we explore an understudied real-world setting of CIL that starts with a strong model pre-trained on a large number of base classes. We hypothesize that a strong base model can provide a good representation for novel classes and incremental learning can be done with small adaptations. We propose a 2-stage training scheme, i) feature augmentation – cloning part of the backbone and fine-tuning it on the novel data, and ii) fusion – combining the base and novel classifiers into a unified classifier. Experiments show that the proposed method significantly outperforms state-of-the-art CIL methods on the large-scale ImageNet dataset (e.g. +10% overall accuracy than the best). We also propose and analyze understudied practical CIL scenarios, such as base-novel overlap with distribution shift. Our proposed method is robust and generalizes to all analyzed CIL settings.

1. Introduction

As deep classifiers become more popular for real-world applications, the need for incrementally learning novel classes (novel data) becomes more prevalent. Training a classifier with both old and novel data is not optimal when old data can become unavailable over time [6, 13, 15, 18, 19, 23, 24, 32]. Fewer old data leads to a high imbalance between the old and novel data, and simply fine-tuning the model causes catastrophic forgetting for the old classes [15].

Class-incremental learning (CIL) methods [5, 6, 8, 11, 23, 29, 34] learn to categorize more and more classes over time. However, they typically start their incremental training with a small number of base classes (e.g. only 50), and add an equally small number of new classes at a time. In many practical scenarios, having a large number of base classes can be a more useful starting point for building an applica-

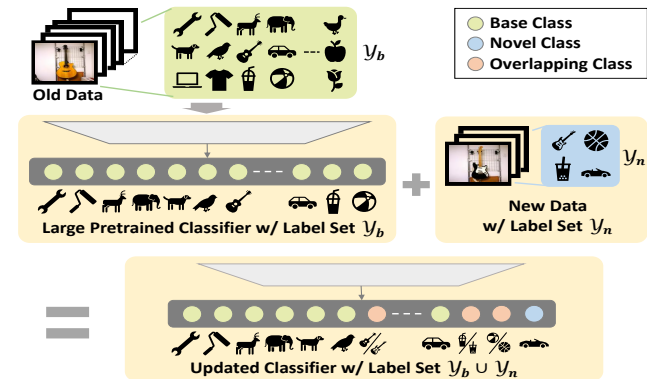


Figure 1. We study the problem of CIL in the setting where there are a large number of base classes. Classes between splits can have overlaps, and data can be sampled from the same or different distributions (e.g. different styles, poses).

tion. For example, a strong model could have been developed to identify different dog breeds, and a small set of additional breeds needs to be added for model update. Moreover, base and novel classes may overlap but may distribute differently, such as a guitar class present in both base and novel classes, with only acoustic guitars in base samples while electric guitars in novel ones.

Some CIL methods use a static model and typically fine-tune the existing parameters with some constraints imposed on parameter changes [3, 13, 32], gradients [17], features [8, 11], or activations [6, 15, 23]. These methods modify the well-trained network weights and risk performance degradation. On the other hand, methods based on dynamic models learn separate parameters for novel tasks, either by expanding the model [24], or introducing a parameter gating function [18, 19]. However, almost all dynamic methods focus on task-incremental learning (TIL). TIL assumes that which task a sample belongs to is known at inference time, and different tasks are inferred individually. This assumption is not realistic if the application needs to distinguish between base and novel classes. A recent work DER [30]

uses a dynamic model for CIL, where it duplicates the entire backbone for novel data and prunes the model.

Current CIL approaches may not be optimal when a large number of base classes (*e.g.* 800) is used to pre-train a strong model. We hypothesize that the well-trained backbone is capable of extracting representative features for the novel data and freezing it partially while learning a small adaptation branch for novel data works better than fine-tuning the whole backbone. We show in a preliminary study that fine-tuning fewer layer blocks outperforms full fine-tuning when using a strong pre-trained model.

Hence, we propose a 2-stage training scheme for CIL starting with a large number of base classes: i) duplicating part of the backbone as the adaptation module and fine-tuning it on the novel data, and ii) combining all the independently trained base and novel classifiers into a unified classifier at each incremental step. Towards this, we propose a score fusion network that enables knowledge transfer between base and novel classes by combining the logits. While the optimal adaptation module size may depend on the old and novel data discrepancy, we show that our score fusion generalizes to different adaptation module sizes.

Most CIL research [6, 8, 11, 23, 29, 34] only consider the scenario where the base and novel label sets are disjoint. Rainbow Memory [5] (RM) is the only exception where label sets are identical among tasks but class frequencies differ. In this work, we explore a more general setting where some novel classes can overlap with base classes, potentially with a different distribution (*e.g.*, different styles, poses), as shown in Figure 1. We show how our score fusion can handle overlapping classes by using a knowledge pooler to combine their base and novel logits.

In summary, we provide three contributions:

- We propose a 2-stage CIL training strategy. In stage-I, instead of tuning the base network and risking catastrophic forgetting, we replicate part of the network and fine-tune the extra branch on novel data. We show that as we start with a strong pre-trained base network, this approach outperforms state-of-the-art CIL methods.
- We propose a new score fusion algorithm for stage-II, where we unify the classifiers for base and novel data into one by consolidating their output logits.
- We generalize CIL to a broader and more challenging scenario: base and novel classes can partially overlap, and the overlap classes may have changed distributions (*e.g.* new style). We show that our method is robust and generalizes to this new scenario.

2. Related work

Prior work in continual learning can be mainly categorized into two streams: methods using a (1) static and (2)

dynamic model architecture. For the first group, the general approach for incremental training is to impose constraints on parameter changes [3, 13, 32], gradients [17], features [8, 11, 25], or activations [2, 6, 10, 15, 23] when fine-tuning the network with novel samples. They can be either memory-free [3, 13, 15, 32] which only rely on the novel data and the original point of convergence in parameters, or be memory-based [2, 6, 8, 11, 13, 17, 23, 29, 34] which relaxes the constraint and keep exemplars of old data for replay. The exemplars can be selected in many ways, such as herding [23], random sampling [10], uncertainty sampling [5].

The second line of work adopt dynamic models, where task-specific parameters are introduced to prevent interference among tasks, which can be achieved by growing the model [14, 16, 21, 22, 24, 26, 30, 31, 33], or introducing a gating function on the parameters [1, 18, 19]. However, most works in this line either focus only on the performance of the target task [16, 21, 22, 26], or assume task label is available at inference time [12, 14, 18, 19, 24, 31, 33]. The exceptions infer the task identity typically by learning a routing classifier [1, 4], or estimating confidence scores [27]. However, they do not produce a unified head for all the classes.

Our proposal is a two-stage method in the offline setting with a dynamic model. In stage-I, we expand the network by cloning and fine-tuning partial parameters from the original backbone. While DER [30] shares some commonalities with our first stage, we differ in stage-II. We show the downside of retraining a linear classifier for all classes solely with the data in the memory as in [5, 30], and propose to combine the logits of the old and new expert classifiers. We also investigate consolidating knowledge of overlapping classes, which is understudied in the literature.

Our method bears superficial similarity to transfer learning and few-shot learning in that it uses a large number of base classes. However, they focus only on the performance of novel classes. Such methods underperform on overall average performance in our experiments.

3. CIL with strong pre-trained models

In this section, we present the problem formulation and the details of the proposed method.

3.1. Problem formulation

Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$, where x_i and y_i are data and label respectively, the goal of a traditional classification network is to learn a feature extractor $h(\cdot; \Phi) \in \mathbb{R}^k$ and a linear classifier $\mathbf{W} \in \mathbb{R}^{k \times |\mathcal{Y}|}$, where \mathcal{Y} is the label set of \mathcal{D} . This is usually obtained by minimizing the cross-entropy loss,

$$L_{ce} = -\frac{1}{N} \sum_{i=1}^N \log \hat{p}^{(y_i)}(\mathbf{x}_i) \quad (1)$$

where

$$\hat{p}(\mathbf{x}) = \sigma(\mathbf{W}^T h(\mathbf{x}; \Phi)), \quad (2)$$

$\sigma(\cdot)$ is the softmax function, and $\mathbf{v}^{(l)}$ is the l^{th} element of \mathbf{v} .

In class-incremental learning (CIL), we first focus on one incremental step, and generalize into multiple steps in Section 3.3. Given a base model \mathcal{M}_b pre-trained on a label set \mathcal{Y}_b using the base dataset \mathcal{D}_b , we augment \mathcal{Y}_b with another label set \mathcal{Y}_n using dataset \mathcal{D}_n , *i.e.*, the new all-labels set $\mathcal{Y}_a = \mathcal{Y}_b \cup \mathcal{Y}_n$. Most prior works [6, 8, 11, 23, 29, 34] focus on the CIL scenario where the label sets are fully disjoint, *i.e.* $\mathcal{Y}_b \cap \mathcal{Y}_n = \emptyset$, while in practice, classes can overlap between base and novel splits [5], where $|\mathcal{Y}_a| < |\mathcal{Y}_b| + |\mathcal{Y}_n|$. A full-overlapping extreme case is $\mathcal{Y}_n \subseteq \mathcal{Y}_b$, which means the new label set $\mathcal{Y}_a = \mathcal{Y}_b$. Despite the overlap, base and novel samples can be very different in each of the overlapping classes, such as having different poses or styles.

An intuitive solution to learn the combined label set is to use $\mathcal{D}_b \cup \mathcal{D}_n$ to train a standard classifier with eq. (1). However, as motivated in the CIL literature, \mathcal{D}_b often becomes unavailable over time in practice. In this case, \mathcal{D}_n dominates the training set and causes catastrophic forgetting, degrading base class performance significantly.

Traditional CIL methods divide existing datasets’ classes evenly into multiple splits, each with a small number of classes (*e.g.* 50). Typically, the initial model is briefly trained only on one of the small splits. Instead, we explore a real-world scenario that is understudied in the literature, where the pre-trained model has been sufficiently trained on a large number of classes (*e.g.* 800 classes). In following incremental steps $|\mathcal{Y}_b| \gg |\mathcal{Y}_n|$, so the smaller novel splits bring comparably little additional information.

3.2. Can minimal network modification work well?

If we freeze the pre-trained network, we avoid catastrophic forgetting. However, this representation might not generalize to the novel classes, especially in traditional CIL where the first-step model is pre-trained with a small number of classes. To mitigate this issue, most prior CIL work [6, 8, 11, 23, 29, 34] update the feature extractor with novel data. Regularization in the loss function (such as penalizing changes in predictions [15, 23] or network weights [13, 32], or re-training on saved exemplars [29, 34]) can reduce catastrophic forgetting to a certain extent, but it still remains prevalent and unsolved.

In this paper, we propose to add extra capacity to the network. We differ from prior network-growing approaches [24, 26] in two ways. First, our goal is CIL rather than TIL, *i.e.*, the network must avoid confusion between novel and base classes. This is not solved by merely freezing existing network weights, which by construct only avoids forgetting in classifying among old classes. More importantly, CIL with a large number of base classes allows

our method to leverage the powerful pre-trained model. We hypothesize that drastic changes in the existing network is unnecessary, and modifying only the top few layers will suffice. Before we introduce our method, we validate this hypothesis and our strategy of minimal network changes with two motivational experiments.

We first analyze *when there are many base classes, how much (or little) benefit is in optimizing representations with the novel data*. Starting from a pre-trained ResNet10 [9] as \mathcal{M}_b , we train a classifier for \mathcal{Y}_n only using new data \mathcal{D}_n , with or without fine-tuning $h(\cdot; \Phi)$. As shown in Figure 2 (left), the gap in novel class accuracy between fine-tuning and fixing the feature extractor is large when the number of base classes is small. The gap significantly reduces with increasing number of base classes, but does not disappear even with $|\mathcal{Y}_b| = 800$. This indicates that fully changing the representations for the novel data is still beneficial for learning \mathcal{D}_n , but with greatly reduced significance.

We then explore *whether we need to optimize the entire pre-trained feature extractor using novel data*. Figure 2 (right) presents the novel class accuracy of fine-tuning smaller subsets of the four convolutional blocks in ResNet10. With weaker pre-trained models (*e.g.* $|\mathcal{Y}_b| = 40, 100, 200$), almost all the layers in the backbone need to be fine-tuned to gain good performance in novel classes. However, for strong pre-trained models (*i.e.*, $|\mathcal{Y}_b| = 800$), fine-tuning anything beyond the last convolutional block (*i.e.*, layer4) degrades the strong performance.

3.3. Training pipeline

Inspired by these observations, we propose a 2-stage incremental training pipeline (Figure 3). We first formulate using one incremental step (base+novel):

Stage-I – Feature augmentation (FA): We split the pre-trained feature extractor $h(\cdot; \Phi)$ into two sub-networks, where the encoder with parameter Φ_s is followed by the encoder with parameter Φ_b , $\Phi = \{\Phi_s, \Phi_b\}$. To optimize the features for \mathcal{D}_n without forgetting the ones for \mathcal{D}_b , we expand the features by cloning the sub-network of Φ_b to the branch Φ_n as the adaptation module, and fine-tuning Φ_n and the weights of the novel class classifier $\mathbf{W}_n \in \mathbb{R}^{k \times |\mathcal{Y}_n|}$ on \mathcal{D}_n with the loss of eq. (1). The shared Φ_s is frozen. This setup ensures no forgetting in the old representation, while enabling feature learning to accommodate new knowledge. While the optimal size of the frozen Φ_s depends on the data discrepancy between base and novel splits, which itself can be a research topic to explore, we adopt the last convolutional block (*i.e.* layer4) as Φ_b as it suffices for most cases, and Φ_s is the layer1-3 blocks.

Stage-II – Fusion: After the first-stage training, we have the base and novel classifiers, $\mathcal{M}_b(\cdot; \Phi_s, \Phi_b, \mathbf{W}_b)$ and $\mathcal{M}_n(\cdot; \Phi_s, \Phi_n, \mathbf{W}_n)$, optimized for \mathcal{Y}_b and \mathcal{Y}_n respectively. We introduce in Section 4 our score fusion scheme to combine the knowledge of the two networks and get a

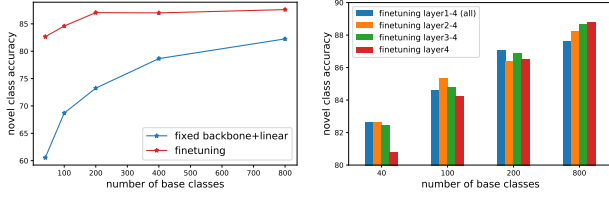


Figure 2. Left: The comparison between fine-tuning and freezing representations. Right: The novel class accuracy of fine-tuning each pre-trained model with different numbers of layers.

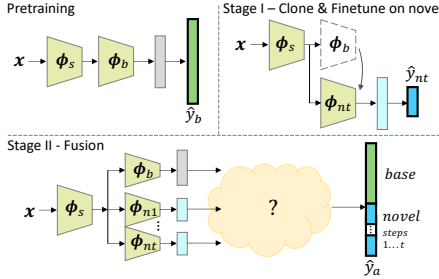


Figure 3. General training pipeline.

unified classifier of the label set \mathcal{Y}_a , where $\mathcal{Y}_a = \mathcal{Y}_b \cup \mathcal{Y}_n$. **Multiple incremental steps:** We maintain one extra branch per new step: $\{\Phi_b, \mathbf{W}_b, \Phi_{n1}, \mathbf{W}_{n1}, \dots, \Phi_{nT}, \mathbf{W}_{nT}\}$. In each novel step $t \in \{1..T\}$, we initialize the new branch Φ_{nt} with Φ_b , fine-tune it, and re-fuse all $t + 1$ branches.

4. Fusion: unifying base and novel classifiers

4.1. Baseline using routing

We analyze fusion baselines with only one incremental step, although they can naturally extend to additional steps. We first explore two intuitive baselines that rely on each classifier’s prediction and a routing function $\hat{r}(\cdot) \in \{0, 1\}$ to decide whether a sample is from \mathcal{D}_b or \mathcal{D}_n . The final prediction \hat{y} is then assigned by \hat{y}_b if $\hat{r}(\mathbf{x}) = 0$ and \hat{y}_n if $\hat{r}(\mathbf{x}) = 1$, where $\hat{y}_d = \arg \max_l \hat{p}^{(l)}(\mathbf{x}; \Phi_s, \Phi_d, \mathbf{W}_d)$, for $d \in \{b, n\}$. We explore two routing functions:

Confidence-based routing (Fig. 4a) uses the confidence score of individual classifiers as the proxy where the routing function is defined as $\hat{r}(\mathbf{x}) = \mathbb{1}_{[Conf_b < Conf_n]}$, where $\mathbb{1}_{[\cdot]}$ is the indicator function, and the confidence score $Conf_d = \max_l \hat{p}^{(l)}(\mathbf{x}; \Phi_s, \Phi_d, \mathbf{W}_d)$, for $d \in \{b, n\}$.

Learning-based routing (Fig. 4b) directly learns a routing classifier with available data. Following prior work in CIL [23], we keep few-shot exemplars \mathcal{E} from all past incremental steps in the memory to supplement the novel data \mathcal{D}_n . The routing classifier is formulated as

$$\hat{r}(\mathbf{x}) = \sigma(\mathbf{W}_r^T (\mathbf{h}_b \oplus \mathbf{h}_n)), \quad (3)$$

where \oplus denotes vector concatenation. For $d \in \{b, n\}$, the feature from each feature extractor branch, $\mathbf{h}_d =$

$h(\mathbf{x}; \Phi_s, \Phi_d)$, is frozen to preserve the information, and the linear weights $\mathbf{W}_r \in \mathbb{R}^{2k \times 2}$ are the only parameters for learning the routing classifier.

The routing loss is a binary cross-entropy loss, *i.e.*

$$l_{rt}(\mathbf{x}, r) = -(1 - r) \log \hat{r}^{(0)}(\mathbf{x}) - r \log \hat{r}^{(1)}(\mathbf{x}), \quad (4)$$

where $r = \mathbb{1}_{[\mathbf{x} \in \mathcal{D}_n]}$ is the split label for \mathbf{x} . However, there is a large base-novel sample imbalance due to $|\mathcal{D}_n| \gg |\mathcal{E}|$. To address this, we re-balance the class losses by re-weighting:

$$L_{rt-bal} = \frac{1}{2|\mathcal{E}|} \sum_{\mathbf{x}_i \in \mathcal{E}} l_{rt}(\mathbf{x}_i, r_i) + \frac{1}{2|\mathcal{D}_n|} \sum_{\mathbf{x}_i \in \mathcal{D}_n} l_{rt}(\mathbf{x}_i, r_i). \quad (5)$$

While routing is a way to get the prediction for all classes in \mathcal{Y}_a , these baselines cannot produce a unified probability distribution for all classes. Also, the routing function’s prediction error will propagate to the final class prediction.

4.2. General score fusion network

Our proposed score fusion network (Figure 4c) integrates the knowledge and output of all branches and generates a unified probability distribution. Modifying the existing well-trained classifiers with few samples leads to overfitting while being susceptible to class imbalance in the limited data, but freezing everything rules out knowledge transfer opportunities between branches. In addition, predictions regarding overlapping classes need unified scores from base and novel classifiers. We introduce knowledge-preserving transfer, overlap knowledge integration schemes, and balanced optimization to address these issues.

Knowledge-preserving transfer. After the stage-I training, we obtain expert models for each step $\{\Phi_d, \mathbf{W}_d\}$ (for $d \in \{b, n1, \dots, nt\}$), where the probability of the individual classifier is computed by applying the softmax function on the logit score $\mathbf{z}_d = \mathbf{W}_d^T \mathbf{h}_d$, $\mathbf{W}_d \in \mathbb{R}^{k \times |\mathcal{Y}_d|}$. To prevent overfitting to the small dataset of $\mathcal{E} \cup \mathcal{D}_{nt}$, we propose to also freeze all classifier weights (*i.e.* $\mathbf{W}_b, \mathbf{W}_{n\tau}$, in addition to $\Phi_s, \Phi_b, \Phi_{n\tau}$, $\tau \in [1, t]$) to preserve their capability of distinguishing classes within the same split. To additionally enable the knowledge transfer between the two splits, we use $\mathcal{E} \cup \mathcal{D}_{nt}$ to learn $\mathbf{W}_{dd'} \in \mathbb{R}^{k \times |\mathcal{Y}_d|}$, $d, d' \in \{b, n1, \dots, nt\}$, $d \neq d'$ (randomly initialized), connecting d' branch’s features to d branch’s logits, which are used to learn the delta logits for knowledge transfer, to be added to the final logits for the d branch:

$$\Delta \mathbf{z}_d = \sum_{d'=b, n1, \dots, nt}^{d' \neq d} \mathbf{W}_{dd'}^T \mathbf{h}_{d'}, \quad (6)$$

$$\tilde{\mathbf{z}}_d = \mathbf{z}_d + \Delta \mathbf{z}_d. \quad (7)$$

Overlap knowledge integration. To get a unified classifier with probability distribution for all the classes, we can

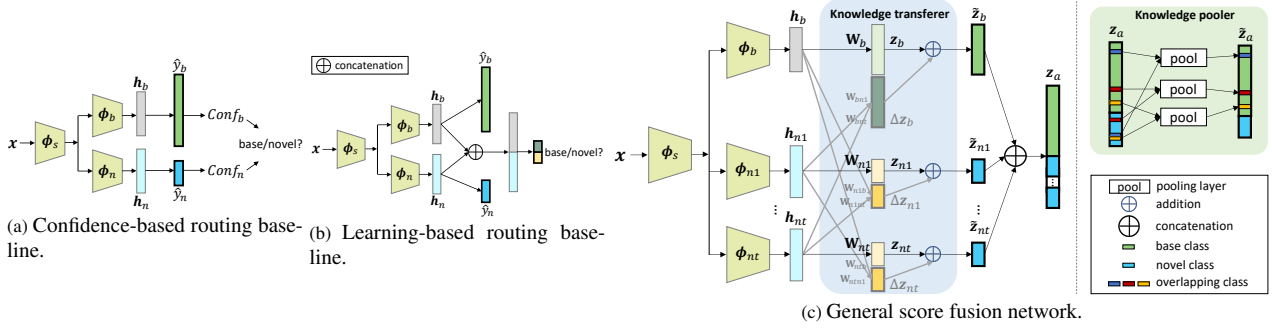


Figure 4. **Fusion methods.** (a-b) Routing as a baseline. (c) Our proposed network for score fusion.

combine the logit scores of the base and novel branches by concatenation, *i.e.* $\mathbf{z}_a = \tilde{\mathbf{z}}_b \oplus \tilde{\mathbf{z}}_{n1} \oplus \dots \oplus \tilde{\mathbf{z}}_{nt}$ where $\tilde{\mathbf{z}}_d \in \mathbb{R}^{|\mathcal{Y}_d|}$, $d \in \{b, n1, \dots, nt\}$, and obtain the posterior probability with $\sigma(\mathbf{z}_a)$. However, when overlapping classes exist ($\mathcal{Y}_d \cap \mathcal{Y}_{d'} \neq \emptyset, \exists d, d'$), they appear in $\tilde{\mathbf{z}}_a$ multiple times. We apply a knowledge pooler to $\mathbf{z}_a \in \mathbb{R}^{\sum_d |\mathcal{Y}_d|}$ to get the final logit $\tilde{\mathbf{z}}_a \in \mathbb{R}^{|\mathcal{Y}_a|}$, as illustrated in Figure 4c (right), that either max-pools or average-pools the multiple logit scores for each overlapping class. Note that $\tilde{\mathbf{z}}_a = \mathbf{z}_a$ when $\mathcal{Y}_d \cap \mathcal{Y}_{d'} = \emptyset, \forall d, d'$. From our experiments, max pooling performs better than average pooling, since the branches do not always simultaneously output high responses to a sample on its class, especially when the data of the same class are very different in base and novel splits.

Balanced optimization. With the final logit score $\tilde{\mathbf{z}}_{a,i}$ for each sample $\mathbf{x}_i \in \mathcal{E} \cup \mathcal{D}_{nt}$, the parameters $\mathbf{W}_{dd'}$ can be optimized with the loss of eq. (1). However, since $|\mathcal{D}_{nt}| \gg |\mathcal{E}|$, the training will be dominated by the novel classes. To balance the probability estimation, we sample a subset $\mathcal{B} \subset \mathcal{E} \cup \mathcal{D}_{nt}$ uniformly over all the classes, where each class in \mathcal{B} has equal number of samples. With this class-balanced sampling, the classification loss becomes

$$L_{cls} = \frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} -\log \hat{p}^{(y_i)}(\mathbf{x}_i) \quad \hat{p}(\mathbf{x}) = \sigma(\tilde{\mathbf{z}}_a). \quad (8)$$

However, in addition to the sample imbalance, our number of classes is also highly imbalanced. Since $|\mathcal{Y}_b| \gg |\mathcal{Y}_{n\tau}|$, base logits have more chances of being the largest than novel logits, so eq. (8) will favor one of the base classes, which may or may not be desired depending on the application. To control base and novel logit balance, we explore two regularization mechanisms.

First, we explicitly train to balance the largest base logit and the largest novel logit scores, by adding a routing auxiliary loss over the maximum score from each split. The routing classifier can be defined as

$$\hat{r}(\mathbf{x}) = \sigma(\mathbf{W}_{r,aux}^T (\max_l \tilde{\mathbf{z}}_b^{(l)} \oplus \max_l \tilde{\mathbf{z}}_{n1}^{(l)} \oplus \dots \oplus \max_l \tilde{\mathbf{z}}_{nt}^{(l)})), \quad (9)$$

where $\mathbf{W}_{r,aux} \in \mathbb{R}^{(t+1) \times (t+1)}$ is the linear routing classifier's weights. The full loss then becomes

$$L_{total} = (1 - \alpha) \cdot L_{cls} + \alpha \cdot L_{rt-bal}, \quad (10)$$

with L_{cls} from eq. (8), L_{rt-bal} from eq. (5) using \hat{r} from eq. (9), and α is a loss weight hyperparameter.

Second, during training we normalize and scale $\mathbf{h}_{d'}$ for base samples by a factor of $\beta \in [0, 1]$ before feeding it into $\mathbf{W}_{bd'}$, $d' \neq b$. This prevents the training from drastically influencing base classes, but limits the knowledge transfer from novel branch features to base classes.

Like all imbalanced learning problems, there is a trade-off between the performance of base and novel classes. Using our two regularization mechanisms provides the flexibility to optimize for the metric that matters more for each customized application. The behavior of these mechanisms will be discussed further in the experiment section.

5. Experiments

In this section, we compare our results with baselines and state-of-the-art CIL methods. We also show ablation studies on the effect of different components and analyze our results in different practical settings. Code will be released on acceptance.

5.1. Experimental setup

Dataset. To study the behavior of strong pre-trained models, we create several data splits from ImageNet [7] to simulate different practical scenarios. For *disjoint-CIL* with disjoint base and novel classes, we first perform our main analysis with one incremental step. Unless otherwise noted, we randomly select 800 base classes and 40 novel classes. Variations are tested in ablation studies. We then test on an existing 10-step asymmetric split scheme with 500 base classes and ten steps of 50 classes each [8], originally designed to test short but numerous incremental steps.

For *overlapping-CIL*, where some or all novel classes overlap with base classes, we experiment with three different base-novel class splits (one incremental step). **1) random class split:** 800 base classes, 40 novel classes with

5 overlapping classes among them. **2) domain-changing split:** ImageNet classes are grouped into two categories, animate and inanimate. Inanimate and animate classes are taken as base and novel classes respectively, with 5 overlapping classes that are randomly selected. **3) style-changing split:** To simulate this scenario with currently available annotations of ImageNet, we select five pairs of classes that are semantically similar to each other, merge each pair into a single class, and use them as five overlapping classes. For example, “pembroke” and “cardigan” can be merged into “corgi”, and “electric guitar” and “acoustic guitar” can be merged into “guitar”. Non-overlapping novel and base classes are chosen randomly.

For overlapping classes in (1) and (2), samples need to be distributed between base and novel splits. We explore splitting each class either **randomly** or by **unsupervised clustering** using K-Means ($K = 2$) on features extracted from a full-ImageNet trained ResNet10 penultimate layer. After splitting all 1000 classes into two splits, we use one split for base classes and the other split for novel classes. That is, the 800 base classes are trained on roughly half of the 800 classes’ data.

Metrics. In prior work, the number of base and novel classes are usually balanced (e.g. $|\mathcal{Y}_b| = |\mathcal{Y}_n| = 50$), so it is natural to simply evaluate using the accuracy over all the classes, i.e. $Acc_{all} = \mathcal{A}(\mathcal{D}_{test})$, where \mathcal{D}_{test} is the testing set and $\mathcal{A}(S) = \frac{1}{|S|} \sum_{(x,y) \in S} \mathbb{1}_{[y = \arg \max_k \hat{p}^{(k)}(x)]}$ either at $t = T$ or averaged over all incremental steps (identical to incremental accuracy). We follow the version implemented in [8] that includes the base step $t = 0$ to facilitate comparison. However, in our setting of strong pre-trained models, $|\mathcal{Y}_b| \gg |\mathcal{Y}_n|$, e.g. 800 and 40. In this case, the overall accuracy will be dominated by the base class performance. A model with high overall accuracy is not guaranteed to perform well for the novel classes. Hence, in addition to the overall accuracy, we also present the accuracy in each split. When there is overlap, we use $Acc_{base} = \mathcal{A}(\{(x,y) : (x,y) \in \mathcal{D}_{test}, y \in \mathcal{Y}_b \setminus (\mathcal{Y}_b \cap \mathcal{Y}_n)\})$, $Acc_{novel} = \mathcal{A}(\{(x,y) : (x,y) \in \mathcal{D}_{test}, y \in \mathcal{Y}_n \setminus (\mathcal{Y}_b \cap \mathcal{Y}_n)\})$ and $Acc_{ovlp} = \mathcal{A}(\{(x,y) : (x,y) \in \mathcal{D}_{test}, y \in \mathcal{Y}_b \cap \mathcal{Y}_n\})$. Note that the classification is still predicted among all classes. An aggregate metric to balance among the per split accuracy is the average of them: $Acc_{avg} = \frac{\sum_{d \in \{b,n\}} Acc_d}{t+1}$ either at the final $t = T$ step or follow the 500+50x10 split implementation [8] by averaging over all steps ($Acc_{avg}^{t=0..10}$). When there is class overlap, we use $Acc_{avg} = (Acc_{base} + Acc_{novel} + Acc_{ovlp})/3$ for partial overlap and $Acc_{avg} = (Acc_{base} + Acc_{ovlp})/2$ for full overlap ($\mathcal{Y}_n \subseteq \mathcal{Y}_b$).

Learning. For stage-II training, 10 exemplars per class are randomly selected to create a class-balanced split, and the experiments are repeated with different random seeds. Hyperparameters α and β are selected with a grid search conducted with the validation set as per the metric of interest.

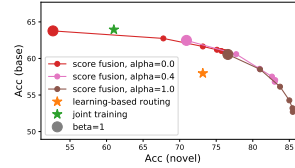


Figure 5. **Design space.** Score fusion performs comparably to joint training and outperforms learning based routing. Optimal α and β values depend on the desired balance between base and novel performance.

For fairness, all compared methods use the same memory constraint and dataset sizes. Dataset subsampling are also identical except where results are from other papers.

5.2. Disjoint-CIL

We first study the most common scenario in the literature, disjoint-CIL, where the set of base and novel classes are disjoint. In the following, we present the comparison to the baselines, along with extensive ablations and analysis of the proposed method.

Design choices. We first study the balancing behavior of the two regularization methods in score fusion, routing loss and feature scaling. For demonstration, we show a hyperparameter sweep in Figure 5 with $\alpha \in \{0, 0.4, 1\}$, $\beta \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$. Actual hyperparameters are selected on validation data instead. As we increase α (routing loss weight) and decrease β (scaling factor), our novel class performance is boosted significantly, while the drop in base class accuracy is comparably minor. The effects of α and β are similar, but changing both gives us the flexibility to reach the full range of base-novel balance operating points.

The optimal hyperparameter selection depends on the specific design choices and desired metric, such as balancing between base and novel (best Acc_{avg}) or prioritizing novel classes (best Acc_{novel}). In the following experiments, we will present the score fusion results optimizing for these practical scenarios: $best-Acc_{all}$, $best-Acc_{avg}$ and $best-balanced$ (optimize $\frac{Acc_{all} + Acc_{avg}}{2}$).

Comparisons to CIL methods. In Table 1 (one novel step) and Table 2 (multiple novel steps), we compare to state-of-the-art class-incremental learning methods and the joint training oracle that uses all base data but unavailable to us. For score fusion, we report the results of three operating points, the one with the best Acc_{all} , the best Acc_{avg} , and the most balanced performance of these two metrics. Fine-tuning and LwF [15] are simple baselines that do not adopt image replay and perform especially badly. iCaRL [23], BiC [29], PODNet [8], and WA [34] keep few-shot exemplars to mitigate forgetting, but these methods still underperform on base classes since they use a small amount of data to modify the well-trained representation layers. We outperform other methods under the multiple novel step

Table 1. Comparison to SOTA class-incremental learning methods. Our method outperforms without additional parameters and pushes the performance further with additional parameters.

| Method | ResNet10 | | | | | ResNet18 | | | | |
|--|-------------|--------------|--------------|---------------|--------------|-------------|--------------|--------------|---------------|--------------|
| | # of params | Acc_{all} | Acc_{base} | Acc_{novel} | Acc_{avg} | # of params | Acc_{all} | Acc_{base} | Acc_{novel} | Acc_{avg} |
| fine-tuning | 4.9M | 4.18 | 0.01 | 87.63 | 43.82 | 11.2M | 4.25 | 0.00 | 89.37 | 44.68 |
| LwF [15] | | 9.50 | 5.54 | 88.53 | 47.04 | | 9.50 | 5.46 | 90.30 | 47.88 |
| iCaRL [23] | | 16.26 | 13.91 | 63.40 | 38.66 | | 10.65 | 8.15 | 60.80 | 34.78 |
| BiC [29] | | 30.30 | 27.55 | 85.20 | 56.38 | | 31.50 | 28.75 | 86.60 | 57.68 |
| WA [34] | | 51.33 | 52.33 | 31.40 | 41.87 | | 54.79 | 55.17 | 47.20 | 51.19 |
| DER w/o P [30] | 9.8M | 52.31 | 52.43 | 50.10 | 51.27 | - | - | - | - | - |
| score fusion (ours) best- Acc_{all} | 8.6M | 63.24 | 63.77 | 52.67 | 58.22 | 19.6M | 69.45 | 70.01 | 58.13 | 64.07 |
| score fusion (ours) best-balanced | | 62.15 | 61.49 | 75.37 | 68.43 | | 67.36 | 66.61 | 82.37 | 74.49 |
| score fusion (ours) best- Acc_{avg} | | 58.90 | 57.73 | 82.40 | 70.06 | | 65.83 | 64.85 | 82.50 | 75.17 |
| score fusion (ours, fc-only) best- Acc_{all} | 4.9M | 62.65 | 63.56 | 44.53 | 54.05 | 11.2M | 68.79 | 69.58 | 53.07 | 61.32 |
| score fusion (ours, fc-only) best-balanced | | 61.01 | 60.81 | 65.07 | 62.94 | | 66.76 | 66.50 | 71.83 | 69.17 |
| score fusion (ours, fc-only) best- Acc_{avg} | | 57.91 | 57.24 | 71.57 | 64.40 | | 65.89 | 65.49 | 73.77 | 69.63 |
| joint learning (oracle) | 4.9M | 63.80 | 63.94 | 61.00 | 62.47 | 11.2M | 70.32 | 70.43 | 68.20 | 69.32 |

Table 2. Multiple novel steps disjoint-CIL with ResNet10, 500 base classes with ten 50-class novel increments, random class split. For ResNet18, BiC and PODNet results are from [8], and we use their experimental setup and 20/class memory constraints.

| Network Method | ResNet10 | | ResNet18 | |
|---------------------------------------|--------------|-----------------------|--------------|-----------------------|
| | inc. acc. | $Acc_{avg}^{t=0..10}$ | inc. acc. | $Acc_{avg}^{t=0..10}$ |
| BiC [29] | - | - | 44.31 | - |
| PODNet [8] | - | - | 64.13 | - |
| score fusion (ours) best- Acc_{all} | 62.70 | 61.45 | 67.48 | 65.66 |
| score fusion (ours) best-balanced | 61.03 | 65.34 | 65.95 | 69.40 |
| score fusion (ours) best- Acc_{avg} | 55.31 | 66.23 | 61.06 | 70.67 |

benchmark as well, including PODNet [8] which proposed the benchmark. DER w/o P [30] is a recent approach that freezes the feature extractor of the base network and creates an identical branch to learn novel features. While similar to the proposed method, the architecture of the additional feature extractor branch and the training objectives are different. In addition, they relearn the final linear classifier with limited data, which we show is less effective in our ablation studies (Table 3). Compared with existing CIL methods with the same backbone, the proposed method achieves significantly better performance in both Acc_{all} and Acc_{avg} and generalize well into multiple novel step scenarios.

For fair comparison, we also present our results branching at the classifier layer (“fc-only”) to bring our number of network parameters on par with most compared methods, i.e. $\Phi_s = \Phi$, $\Phi_b = \Phi_n = \emptyset$, and $\mathbf{h}_b = \mathbf{h}_n$. Our results still outperform all compared methods, and interestingly, our ResNet10 results outperform others’ ResNet18 results despite using fewer parameters.

Comparisons to fusion baselines. In Table 3, we explore two lines of fusion methods in stage-II training, routing, and score fusion. We show that our observations hold for different backbones (ResNet18/ResNet50) and more novel classes in supplemental.

Among routing methods, confidence-based routing has

Table 3. Disjoint-CIL analysis with ResNet10, 40 novel classes using random class split. (Results for ResNet18/ResNet50 architectures and for 200 novel classes are in supplemental)

| Method | Acc_{all} | Acc_{base} | Acc_{novel} | Acc_{avg} |
|--|--------------|--------------|---------------|--------------|
| confidence-based routing | 41.58 | 39.26 | 88.00 | 63.63 |
| learning-based routing w/ L_{rt-bal} | 58.69 | 57.97 | 73.17 | 65.57 |
| oracle routing w/ L_{rt-bal} | 58.57 | 57.48 | 80.50 | 68.99 |
| FeatCat+RT | 55.94 | 56.08 | 53.27 | 54.67 |
| LogitCat+RT | 58.36 | 58.56 | 54.30 | 56.43 |
| LogitCat+FT | 59.39 | 59.34 | 60.57 | 59.95 |
| FA (ours) + BiC [29] | 63.21 | 63.66 | 54.20 | 58.93 |
| FA (ours) + WA [34] | 62.77 | 64.19 | 34.30 | 49.25 |
| score fusion (ours) best- Acc_{all} | 63.24 | 63.77 | 52.67 | 58.22 |
| score fusion (ours) best-balanced | 62.15 | 61.49 | 75.37 | 68.43 |
| score fusion (ours) best- Acc_{avg} | 58.90 | 57.73 | 82.40 | 70.06 |
| joint learning (oracle) | 63.80 | 63.94 | 61.00 | 62.47 |

the worst Acc_{base} , and around half of the base samples are misclassified as novel. Learning-based routing improves the result for Acc_{base} , and reaches a more balanced performance. The oracle routing performance (trained on all base and novel data rather than 10-shot) is also shown in the table for reference, and it mainly improves in novel classes and the gap in Acc_{base} remains.

We then compare the score fusion method with a set of baselines in the proposed framework in Table 3. *FeatCat* + *RT* retraining a linear classifier with features \mathbf{h}_b and \mathbf{h}_n concatenated. *LogitCat* concatenates logits \mathbf{z}_{base} and \mathbf{z}_{novel} , and retraining (*RT*) or fine-tunes (*FT*) the linear weights \mathbf{W}_b and \mathbf{W}_n . Since the original linear weights were already trained using a large amount of data, retraining them from scratch on 10-shot with feature or logit concatenation leads to inferior results than fine-tuning the linear weights, as shown in the table. This supports our idea of preserving the weights of the original linear classifiers.

Next, we apply two score fusion baselines, BiC [29] and WA [34], to our network (FA) after stage-I training. Both

Table 4. Disjoint-CIL with ResNet10, 40 novel classes using the inanimate-animate split (wider base-novel gap). Score fusion can generalize to fine-tuning more/less layers in stage-I if necessary.

| Trainable | Method | Acc_{all} | Acc_{base} | Acc_{novel} | Acc_{avg} |
|---------------------|-----------------------------------|-------------|--------------|---------------|-------------|
| layer3,4 | score fusion (ours) best-balanced | 57.08 | 55.27 | 83.10 | 69.18 |
| layer4 | score fusion (ours) best-balanced | 57.12 | 55.50 | 80.37 | 67.93 |
| layer4 (conv2 only) | score fusion (ours) best-balanced | 57.15 | 55.73 | 77.37 | 66.55 |

methods estimate a scalar to balance between the base and novel logits, and they estimate it by validation data and weight norms respectively. Using these methods with our network preserves the base class performance better than *RT* or *FT* methods, and FA+BiC on ResNet10 even performs similarly to ours (but not on ResNet18). Although we note that using BiC and WA directly performs poorly as shown in Table 1.

Robustness. To test the robustness of the proposed method, we study score fusion with a larger number of novel classes and deeper networks. Results of 800 base and 200 novel classes instead of 40 and the results for different architectures are in supplemental. The same trends in Table 3 hold. **Generalization.** When the base and novel data are very different, *e.g.* in Table 4 where we use the domain-changing split (inanimate as base classes and animate for novel), fine-tuning more layers in the first stage may produce higher accuracy on the novel data. We show that our score fusion generalizes to different ways of branching the network.

5.3. Overlapping-CIL

In practice, base and novel classes may not be mutually exclusive. In this section, we study the overlapping-CIL scenario where a subset of base and novel classes overlap. As a recap of Section 5.1, we analyze five different ways to split base and novel classes and samples within each class. Unless otherwise noted, we test on ResNet10 with 800 base, 40 novel with 5 overlapping classes.

Table 5 shows the result when the style of the overlapping classes changes from base to novel. We simulate this by merging pairs of similar classes that are siblings in the WordNet hierarchy [20, 28], and placing the two classes in each pair into the base and novel split respectively. Compared to average pooling, applying max pooling to the logits of overlapping classes performs better in Acc_{ovlp} with little drawback in Acc_{base} . Hence, in our score fusion results, we adopt max pooling to fuse the logits. Confidence-based routing performs the worst for the base classes. Our score fusion achieves the best performance for both splits even when the samples of overlapping classes are in a different style from base to novel.

Due to space constraints, we show in the supplemental results for the random base-novel-overlap split scenario and the domain change (inanimate and animate, respectively, and 5 inanimate classes as overlapping classes) scenario. The conclusions are identical and our score fusion general-

Table 5. Overlapping-CIL results of style-changing splits (merge similar classes into one). See supplemental for overlapping-CIL with random and domain-changing split.

| Method | Acc_{all} | Acc_{base} | Acc_{novel} | Acc_{ovlp} | Acc_{avg} |
|--|--------------|--------------|---------------|--------------|--------------|
| confidence-based routing | 40.70 | 38.17 | 86.97 | 80.00 | 68.38 |
| learning-based routing w/ L_{rt-bal} | 59.14 | 58.63 | 72.91 | 51.73 | 61.09 |
| oracle routing w/ L_{rt-bal} | 59.57 | 58.83 | 78.63 | 51.20 | 62.89 |
| logit concatenation (avg pool) | 63.39 | 63.63 | 64.57 | 40.40 | 56.20 |
| logit concatenation (max pool) | 63.50 | 63.38 | 64.46 | 69.60 | 65.81 |
| score fusion (ours) best- Acc_{all} | 63.81 | 64.35 | 53.76 | 56.13 | 58.08 |
| score fusion (ours) best-balanced | 61.02 | 60.02 | 79.58 | 75.60 | 71.73 |
| score fusion (ours) best- Acc_{avg} | 57.58 | 56.18 | 83.77 | 77.73 | 72.56 |
| joint learning (oracle) | 64.33 | 64.28 | 62.06 | 76.80 | 67.71 |

Table 6. Fully-overlapping-CIL with all 40 novel classes overlapping with the 800 base classes, *i.e.* $\mathcal{Y}_n \subset \mathcal{Y}_b$.

| Backbone | Method | Acc_{all} | Acc_{base} | Acc_{ovlp} | Acc_{avg} |
|----------|--|--------------|--------------|--------------|--------------|
| ResNet10 | pre-trained model | 57.69 | 58.03 | 51.20 | 54.62 |
| | confidence-based routing | 35.43 | 33.01 | 81.50 | 57.25 |
| | learning-based routing w/ L_{rt-bal} | 52.82 | 52.12 | 66.10 | 59.11 |
| | fine-tuning | 3.86 | 0.00 | 77.20 | 38.6 |
| | iCaRL [23] | 20.23 | 19.77 | 28.90 | 24.34 |
| | RM [5] | 21.10 | 18.12 | 77.70 | 47.91 |
| | score fusion (ours) best- Acc_{all} | 57.82 | 57.80 | 58.23 | 58.02 |
| | score fusion (ours) best-balanced | 55.66 | 54.81 | 71.90 | 63.36 |
| | score fusion (ours) best- Acc_{avg} | 53.42 | 52.28 | 75.13 | 63.70 |
| | joint learning (oracle) | 47.63 | 47.32 | 53.60 | 50.46 |
| ResNet18 | iCaRL [23] | 18.12 | 17.04 | 38.60 | 27.82 |
| | RM [5] | 21.89 | 28.78 | 80.90 | 49.84 |

izes to both scenarios, with overlapping class samples either randomly split or clustered.

Additionally, Table 6 shows an extreme full-overlap scenario where all novel classes are overlapping with base classes (random base-novel split, split overlapping classes by clustering). This simulates the practical scenario where one needs to adapt a portion of existing classes to a new domain. We compare to prior incremental learning work. These methods underperform due to changing the backbone network weights. Even Rainbow Memory [5] which explicitly tackles fully overlapping classes, underperforms, partially because it assumes the distribution of each overlapping class does not change. In comparison, our method generalizes well to this scenario.

6. Conclusion and Limitations

In this work, we investigated CIL in the context of a pre-trained model with large number of base classes. We showed how branching can be an effective solution for learning novel data when using a strong pre-trained model and how it can preserve the learning with the old data. Furthermore, we discuss a novel score fusion method that uses both feature and classifier information from old and novel networks and generates a unified classifier. This approach leads to state-of-the-art results for CIL with large number of base classes. Our method can further be improved by using distillation approach to reduce memory footprint.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3931–3940, 2020. [2](#)
- [2] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. Ss-il: Separated softmax for incremental learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 844–853, 2021. [2](#)
- [3] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. [1](#), [2](#)
- [4] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3366–3375, 2017. [2](#)
- [5] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8218–8227, 2021. [1](#), [2](#), [3](#), [8](#)
- [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](#), [2](#), [3](#)
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. [5](#)
- [8] Arthur Douillard and et al. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *ECCV 2020*. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [3](#)
- [10] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Lifelong learning via progressive distillation and retrospection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 437–452, 2018. [2](#)
- [11] 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. [1](#), [2](#), [3](#)
- [12] Steven CY Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. *arXiv preprint arXiv:1910.06562*, 2019. [2](#)
- [13] 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](#), [3](#)
- [14] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In *International Conference on Machine Learning*, pages 3925–3934. PMLR, 2019. [2](#)
- [15] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [1](#), [2](#), [3](#), [6](#), [7](#)
- [16] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pages 97–105. PMLR, 2015. [2](#)
- [17] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476, 2017. [1](#), [2](#)
- [18] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018. [1](#), [2](#)
- [19] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. [1](#), [2](#)
- [20] George A. Miller. Wordnet: A lexical database for english. *Communications of the ACM*, 38:39–41, 11 1995. [8](#)
- [21] Pedro Morgado and Nuno Vasconcelos. Nettare: Tuning the architecture, not just the weights. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3044–3054, 2019. [2](#)
- [22] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *arXiv preprint arXiv:1705.08045*, 2017. [2](#)
- [23] 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](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [24] 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](#), [3](#)
- [25] Xiaoyu Tao, Xinyuan Chang, Xiaopeng Hong, Xing Wei, and Yihong Gong. Topology-preserving class-incremental learning. In *European Conference on Computer Vision*, pages 254–270. Springer, 2020. [2](#)
- [26] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Growing a brain: Fine-tuning by increasing model capacity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2471–2480, 2017. [2](#), [3](#)
- [27] Mitchell Wortsman, Vivek Ramanan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *arXiv preprint arXiv:2006.14769*, 2020. [2](#)

- [28] Tz-Ying Wu, Pedro Morgado, Pei Wang, Chih-Hui Ho, and Nuno Vasconcelos. Solving long-tailed recognition with deep realistic taxonomic classifier. In *European Conference on Computer Vision (ECCV)*, 2020. 8
- [29] Yue Wu and et al. Large scale incremental learning. In *CVPR 2019*. 1, 2, 3, 6, 7
- [30] 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, 7
- [31] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *International Conference on Learning Representations (ICLR)*, 2018. 2
- [32] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 1, 2, 3
- [33] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: A baseline for network adaptation via additive side networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 698–714. Springer, 2020. 2
- [34] 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. 1, 2, 3, 6, 7