

On the Stability-Plasticity Dilemma of Class-Incremental Learning

Dongwan Kim¹

Bohyung Han^{1,2}

Computer Vision Laboratory, ECE¹ & IPAI², Seoul National University

{dongwan123, bhhan}@snu.ac.kr

Abstract

A primary goal of class-incremental learning is to strike a balance between stability and plasticity, where models should be both stable enough to retain knowledge learned from previously seen classes, and plastic enough to learn concepts from new classes. While previous works demonstrate strong performance on class-incremental benchmarks, it is not clear whether their success comes from the models being stable, plastic, or a mixture of both. This paper aims to shed light on how effectively recent class-incremental learning algorithms address the stability-plasticity trade-off. We establish analytical tools that measure the stability and plasticity of feature representations, and employ such tools to investigate models trained with various algorithms on large-scale class-incremental benchmarks. Surprisingly, we find that the majority of class-incremental learning algorithms heavily favor stability over plasticity, to the extent that the feature extractor of a model trained on the initial set of classes is no less effective than that of the final incremental model. Our observations not only inspire two simple algorithms that highlight the importance of feature representation analysis, but also suggest that class-incremental learning approaches, in general, should strive for better feature representation learning.

1. Introduction

Despite the unprecedented success of deep learning [19, 23, 27, 30], most deep neural networks have static use cases. However, real-world problems often require adaptivity to incoming data [17], changes in training environments, and domain shifts [2, 7, 15]. Thus, researchers have been actively working on model adaptation techniques, and have proposed various continual learning approaches so far.

A naïve approach for continual learning is to simply fine-tune a model. However, such a solution is rather ineffective due to a phenomenon known as *catastrophic forgetting* [13], which arises as a result of high *plasticity* of neural networks, *i.e.* parameters important for the old tasks are updated to better fit the new data. On the flip side, enforcing

model *stability* introduces its own set of limitations, mainly the lack of adaptivity to new data. Thus, we encounter the *stability-plasticity dilemma*: how can we balance stability and plasticity such that the model is able to learn new concepts while retaining old ones? Finding an optimal balance between these two opposing forces is a core challenge of continual learning research, and has been the main focus of many previous works [6, 12, 16, 20, 21, 26, 31].

We conduct an in-depth study of recent works in continual learning, specifically concentrating on class-incremental learning (CIL)—a subfield of continual learning—where new sets of classes arrive in an online fashion. We are motivated by the lack of systematic analyses in the field of CIL, which hampers our understanding of how effectively the existing algorithms balance stability and plasticity. Moreover, works that do perform analyses usually focus on the classifier, *e.g.*, classifier bias [1, 10], rather than the intermediate feature representations. However, investigating the stability and plasticity in the feature level is just as important, if not more, because utilizing the model’s full capacity to learn robust representations is critical to maximizing the potential of CIL algorithms.

To measure plasticity, we retrain the classification layer of CIL models at various incremental stages and study how effectively their feature extractors have learnt new concepts. We then investigate stability by measuring feature similarity with Centered Kernel Alignment (CKA) [3, 14] and by visualizing the feature distribution shift using t-SNE [29]. Surprisingly, and possibly concerning, our analyses show that the majority of CIL models accumulate little new knowledge in their feature representations across incremental stages. In fact, most of the analyzed CIL algorithms seem to alleviate catastrophic forgetting by heavily overlooking model plasticity in favor of high stability. Finally, we introduce two simple algorithms based on our observations. The first is an extension of Dynamically Expandable Representations (DER) [31], which demonstrates how our analyses may be used to improve the efficiency and efficacy of CIL algorithms. The second is an exploitative method, which can be interpreted as an extreme case of DER [31] in terms of architectural design; this method shares a sim-

ilar motivation with GDumb [22] in the sense that it raises significant concerns regarding the current state of CIL research. We summarize our contributions below:

- We design and conduct analytical experiments to better understand the balance between stability and plasticity in the feature representations of modern continually learned models.
- We discover that the feature representations of most CIL models are only subject to trivial updates as the model is trained on incremental data. This is a direct result of overweighing the importance of stability over plasticity, and thus, implies a failure of their balance.
- We present two simple but effective CIL algorithms inspired by the results of our analyses. One is an exploit that highlights a major flaw in the current state of CIL research while the other is a variation of an existing method designed for improving performance.

2. Preliminaries

To set the stage for our paper, we first describe the task setting, notations, and related works.

2.1. Task setting and notations

In continual learning, a neural network model is trained on data that arrives incrementally. Formally, after first training a model with an initial dataset \mathcal{D}_0 , additional datasets $\{\mathcal{D}_i\}_{i=1}^N$ arrive in N sequential steps to further update the model. We collectively denote all incremental datasets as $\{\mathcal{D}_i\}_{i=0}^N$ for an N -step setting. In class-incremental learning (CIL), \mathcal{D}_i consists of examples in a set of classes, \mathcal{C}_i , and collectively, $\{\mathcal{C}_i\}_{i=0}^N$, where all classes are unique such that $|\cup_{i=0}^N \mathcal{C}_i| = \sum_{i=0}^N |\mathcal{C}_i|$. For convenience, we refer to the entire dataset as \mathcal{D} , and all classes as \mathcal{C} . Furthermore, we note that \mathcal{D} may refer to either the training dataset or the validation dataset, depending on the context.

We denote the model trained on \mathcal{D}_0 as \mathcal{M}_0 , and, by extension, define the set of all models trained incrementally on $\{\mathcal{D}_i\}_{i=1}^N$ as $\{\mathcal{M}_i\}_{i=0}^N$. Note that \mathcal{M}_i ($i > 0$) is first initialized with the parameters of \mathcal{M}_{i-1} , and trained on \mathcal{D}_i , and optionally with a small exemplar set sampled from $\{\mathcal{D}_i\}_{i=0}^{i-1}$. We assume the model architectures are based on a convolutional neural network, which, at any given stage i , is composed of a feature extractor \mathcal{F}_i , and classifier \mathcal{G}_i as

$$\mathcal{M}_i = \mathcal{G}_i \circ \mathcal{F}_i. \quad (1)$$

The classifier typically refers to a single linear layer, which may be replaced by a cosine classifier in some algorithms¹.

¹The details of the cosine classifier are described in Section 4 of the Appendix.

Experimental setting To better analyze feature representations in CIL settings, we conduct all experiments on a large-scale dataset, ImageNet-1K [28], with a ResNet-18 [9] architecture, which is a widely-adopted architecture for ImageNet-1K experiments. There exist three common settings for class-incremental learning on ImageNet-1K:

1. **B500-5step**: $|\mathcal{C}_0| = 500$, $N = 5$, and $|\mathcal{C}_{i>0}| = 100$.
2. **B500-10step**: $|\mathcal{C}_0| = 500$, $N = 10$, and $|\mathcal{C}_{i>0}| = 50$.
3. **B0-10step**: $|\mathcal{C}_0| = 100$, $N = 9$, and $|\mathcal{C}_{i>0}| = 100$.

The first and second settings are both “pre-trained” settings, where the model is initially trained on 500 classes, then incrementally updated with fewer classes in later stages. This setting is quite challenging in terms of stability since the model must retain the pre-trained knowledge when learning from new data. Moreover, it presents an even greater challenge for plasticity since the model can already extract reasonable feature representations, and thus, must *actively* learn new representations. We elaborate on this in Section 6. Due to space constraints, we present analyses with the B500-5step setting in our main paper, and the latter two in the Appendix. For all settings, 20 exemplars of each previously seen class is stored in the memory for subsequent stages.

2.2. Overview of compared methods

We briefly summarize the main idea of the existing CIL algorithms analyzed in this work.

Naive The naive method adopts simple fine-tuning, where \mathcal{M}_{i+1} is initialized by a fully-trained \mathcal{M}_i . Only the cross-entropy loss is used to train the network, and exemplars are sampled randomly from each previously observed class.

iCARL [26] Incremental Classifier and Representation Learning (iCARL) employs a simple distillation loss alongside the cross-entropy loss. iCARL also proposes herding for exemplar selection, and discovers that nearest-mean-of-exemplars classification can be beneficial for CIL.

LUCIR [10] Learning a Unified Classifier Incrementally via Rebalancing (LUCIR) proposes to use the cosine classifier for feature rebalancing purposes, and alleviates the adverse effects of classifier imbalance by using the cosine between features of the student and teacher models.

SSIL [1] Separated Softmax for Incremental Learning (SSIL) identifies that score bias may be caused by data imbalance, and trains the model with a separated softmax output layer alongside task-wise knowledge distillation.

AANet [16] Adaptive Aggregation Networks (AANet) employs a two branch residual block, where one corresponds to a stable (fixed) block, while the other corresponds to a plastic block. Existing algorithms can be applied to the AANet architecture. We focus on AANet + LUCIR, and for simplicity, we denote AANet + LUCIR as AANet.




Figure 1. Accuracy on the ImageNet-1K validation set after fine-tuning the classification layer of each incremental model (B500-5step setting) with the full ImageNet-1K training data. The black line indicates an Oracle model trained on $\{500, 600, \dots, 1000\}$ classes, and serves as a point of reference for the performance on non-incremental settings.

POD [6] Pooled Outputs Distillation (POD) employs various types of pooling dimensions for knowledge distillation to enforce constraints on feature representations between old and new models.

AFC [12] Adaptive Feature Consolidation (AFC) first estimates the importance of each channel in the feature map based on the expected increase in loss, and then adaptively restricts the updates to the important channels while leaving non-important channels relatively unconstrained.

DER [31] Dynamical Expandable Representations (DER) adds a new feature extractor at each incremental stage, and leaves feature extractors trained on older data fixed while the new feature extractor is updated. For any stage i , the outputs of all $i+1$ feature extractors are concatenated before passing through the classification layer.

3. Re-evaluating Feature Representations

The lack of model stability and/or plasticity often leads to weak performance, where a suboptimal feature extractor is unable to extract meaningful information. Likewise, a suboptimal classifier (*e.g.*, due to classifier bias) further exacerbates this issue. While previous works have extensively studied classifier bias [1, 10, 33], the effects of unbalanced stability and plasticity in the feature extractor has been relatively less explored, and will be the focus of this section.

3.1. Finetuning the classifier on full data

We begin by examining how the performance of feature extractors transforms over incremental steps. To eliminate

any negative effects of an incrementally trained classifier, we freeze the feature extractors $\{\mathcal{F}_0, \dots, \mathcal{F}_N\}$ and train a new classifier for each of the feature extractors on the full ImageNet-1K training data, \mathcal{D} . In essence, we assume that the classifier is optimal (well-fitted to any given feature extractor), and evaluate the strength of a feature extractor by using the accuracy on the ImageNet-1K validation set as a proxy measure. Although retraining the classifier on \mathcal{D} is a breach of CIL protocol, we emphasize that the goal of this experiment is purely to analyze CIL models from the perspective of feature representations.

From here on out, $\mathcal{M}'_j \equiv \mathcal{G}' \circ \mathcal{F}_j$ denotes the combination of the feature extractor from stage j and the retrained classifier, \mathcal{G}' . In other words, \mathcal{F}_j is trained incrementally on $\{\mathcal{D}_i\}_{i=0}^j$, while \mathcal{G}' replaces the original classifier \mathcal{G}_j and is retrained on the entire dataset \mathcal{D} with a frozen \mathcal{F}_j . Note that, while \mathcal{G}_j outputs logits for $\sum_{i=0}^j |\mathcal{C}_i|$ classes, \mathcal{G}' outputs $|\mathcal{C}|$ -dimensional logits. We then define an accuracy metric, $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$, as the accuracy of model \mathcal{M}'_j on the validation dataset at the i^{th} stage, \mathcal{D}_i . Without loss of generality, $\text{Acc}(\mathcal{M}'_j, \mathcal{D})$ denotes the accuracy of model \mathcal{M}'_j on the full ImageNet-1K validation dataset, \mathcal{D} .

3.2. ImageNet-1K accuracies

Figure 1 illustrates the full validation accuracy on ImageNet-1K, $\text{Acc}(\mathcal{M}'_j, \mathcal{D})$, for all the compared algorithms. Each subplot visualizes the accuracy progression for the specified CIL algorithm as well as the Oracle model, \mathcal{M}_j^* , whose feature extractor is trained on $\cup_{i=0}^j \mathcal{D}_i$ all at once, before the classifier is retrained on \mathcal{D} . In essence, the

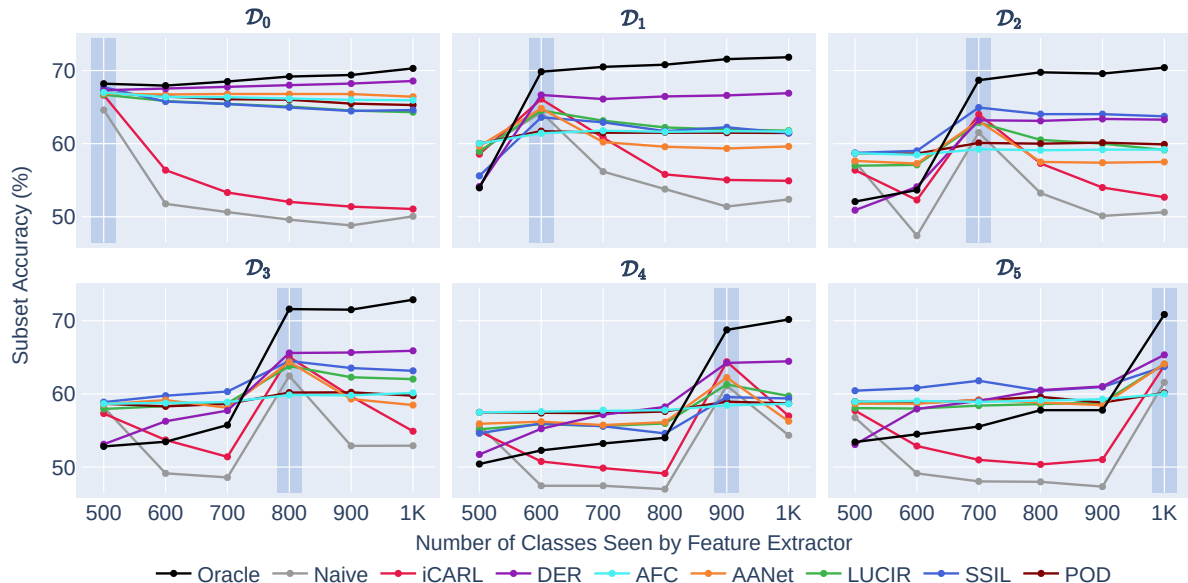


Figure 2. B500-5step subset accuracies. For the sake of visibility, we leave out SSIL and POD from these plots. We highlight the region for model \mathcal{M}'_j in plot \mathcal{C}_i , where $i = j$.

feature extractor of \mathcal{M}'_j represents how an ideal incremental model would perform if important features from previous tasks are not forgotten and new concepts are well learned.

Figure 1 generally exhibit three distinct trends. First, we observe the plots for the Naive and iCARL models, and notice that the accuracy significantly declines after the initial stage, *i.e.*, $\text{Acc}(\mathcal{M}'_0, \mathcal{D}) > \text{Acc}(\mathcal{M}'_j, \mathcal{D}), \forall j > 0$. These results imply that both models are subject to severe catastrophic forgetting, and their ability to extract useful features deteriorates significantly from the first incremental stage. Next, we investigate the subplots for LUCIR, POD, SSIL, AANet, and AFC. Surprisingly, for these five distinct CIL algorithms, the accuracy remains almost unchanged across all incremental stages, despite a few minor variations. Unlike the Oracle model, which shows an almost linear increase in accuracy at each incremental step, these five CIL algorithms maintain the same accuracy, *i.e.*, $\text{Acc}(\mathcal{M}'_0, \mathcal{D}) \approx \text{Acc}(\mathcal{M}'_j, \mathcal{D}), \forall j > 0$. Thus, for these 5 algorithms, the feature extractor of \mathcal{M}'_5 is not particularly stronger than the feature extractor of $\mathcal{M}'_j, \forall j < 5$. Lastly, DER exhibits increasing accuracy with each incremental stage, indicating that the feature extractor does indeed learn new features.

3.3. ImageNet-1K subset accuracies

Diving deeper, we investigate how the accuracy of each subset, \mathcal{D}_i , changes at each incremental stage. Figure 2 illustrates the change in $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i), \forall j$ for each i . For the sake of better visibility, we omit POD and SSIL in this figure and present the full plots in the Appendix.

First, we focus on the gray curve, representing the Naive model. We expected the Naive model to suffer from catas-

trophic forgetting, and overfit to the most recently seen set of classes. Indeed, the results corroborate our intuition; in the first plot on \mathcal{D}_0 , $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_0), \forall j > 0$, performs significantly worse than $\text{Acc}(\mathcal{M}'_0, \mathcal{D}_0)$. Furthermore, we notice that $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$ peaks when $j = i$. Then, $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$ drops off again when $j > i$. These observations all lead to the same conclusion that the naive model suffers from catastrophic forgetting due to its high plasticity. A similar pattern is observed from the models trained by iCARL.

Next, we shift our focus to the black curve corresponding to the Oracle model. We notice that for \mathcal{D}_0 , $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_0)$ does not drop, but rather increases as more classes are added, which suggests that knowledge from $\mathcal{C}_i, \forall i > 0$, can in fact be beneficial for performance on \mathcal{C}_0 . Moreover, $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$ significantly increases when $j = i$. Altogether, the trends exhibited by the Oracle model represent what an ideal CIL model would demonstrate.

Finally, we look at the cyan, orange, and green curves, respectively representing AFC, AANet, and LUCIR. For AFC, $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$ is mostly unchanged $\forall i, j$, suggesting that their features are mostly static across all incremental stages. While this implies that no forgetting occurs, it appears to come at the cost of learning little to no new concepts. Meanwhile, both AANet and LUCIR are mostly stable, but also slightly more plastic than AFC and POD; they quickly forget newly learned concepts since $\text{Acc}(\mathcal{M}'_j, \mathcal{D}_i)$ peaks when $j = i$ but drops back down when $j > i$.

4. Are Incrementally Learned Features Static?

Our analysis in Section 3 suggests that a majority of the compared CIL algorithms appear to have high feature sta-

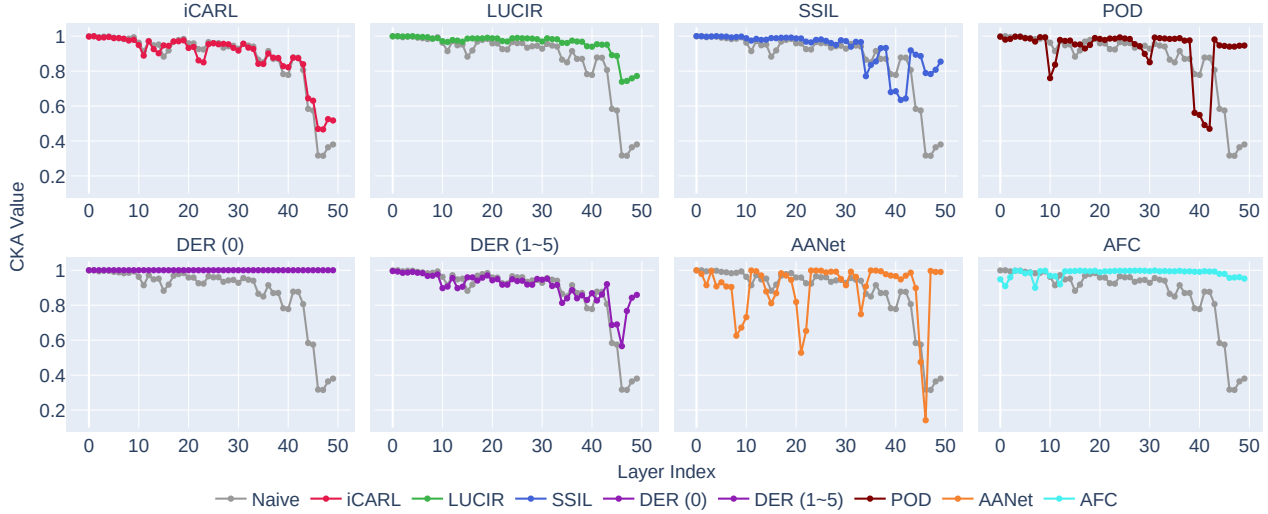


Figure 3. Same-layer CKA values between \mathcal{F}_0 and \mathcal{F}_5 for incremental models trained with each CIL algorithm on the ImageNet B500 5-step setting. The x-axis spans the layer index of ResNet-18, while the y-axis represents CKA. CKA is evaluated using the \mathcal{D}_0 validation set. Each plot is accompanied by the CKA for the naive model, which acts as a point of reference.

bility at the cost of low plasticity. This raises a question: do feature representations remain static across incremental models? In hopes to shed light on this issue, we measure the similarity of intermediate activations and visualize the feature distribution shifts between incremental models.

4.1. Centered Kernel Alignment (CKA)

To analyze intermediate representations of a neural network, we employ Centered Kernel Alignment (CKA) [3, 14], which enables us to quantify the similarity between pairs of neural network representations. CKA has been used to study the effects of increasing depth and width in a network [18], and to understand how the representations of Vision Transformers [5] differ from those of convolutional neural networks [25].

Let us consider two arbitrary layers of a neural network with z_1 and z_2 output features. Given the same set of b inputs, we denote the activation matrices as $\mathbf{X} \in \mathbb{R}^{b \times z_1}$ and $\mathbf{Y} \in \mathbb{R}^{b \times z_2}$. The $b \times b$ Gram matrices $\mathbf{K} = \mathbf{X}\mathbf{X}^T$ and $\mathbf{L} = \mathbf{Y}\mathbf{Y}^T$ are first centered to obtain \mathbf{K}' and \mathbf{L}' , which are then used to compute the Hilbert-Schmidt Independence Criterion (HSIC) [8] as

$$\text{HSIC}(\mathbf{K}, \mathbf{L}) = \frac{\text{vec}(\mathbf{K}') \cdot \text{vec}(\mathbf{L}')}{(b-1)^2}, \quad (2)$$

where $\text{vec}(\cdot)$ denotes the vectorization operation. Finally, CKA normalizes HSIC as follows:

$$\text{CKA}(\mathbf{X}, \mathbf{Y}) = \frac{\text{HSIC}(\mathbf{X}\mathbf{X}^T, \mathbf{Y}\mathbf{Y}^T)}{\sqrt{\text{HSIC}(\mathbf{X}\mathbf{X}^T, \mathbf{X}\mathbf{X}^T)\text{HSIC}(\mathbf{Y}\mathbf{Y}^T, \mathbf{Y}\mathbf{Y}^T)}}. \quad (3)$$

As shown above, CKA is a normalized measure of how similar the $b \times b$ Gram matrices \mathbf{K} and \mathbf{L} are. Given that the Gram matrices themselves reflect the feature relationships among pairs of samples, CKA can be interpreted as a similarity of relationships among the features in \mathbf{X} and \mathbf{Y} .

In terms of comparing feature representations, we highlight three properties that make CKA stand out. First, CKA is invariant to permutations in the columns of \mathbf{X} and \mathbf{Y} , *i.e.*, $\text{CKA}(\mathbf{X}, \mathbf{Y}) = \text{CKA}(\mathbf{X}\mathbf{P}, \mathbf{Y})$, where $\mathbf{P} \in \{0, 1\}^{z_1 \times z_1}$ is an arbitrary permutation matrix. Second, it is invariant to isotropic scaling of \mathbf{X} and \mathbf{Y} . Finally, it can be used to compare activations of layers with different output feature dimensions, *e.g.*, different layers of the same network, or even layers across different architectures. Such advantages make CKA suitable for analyzing the feature representations of class-incremental models.

4.2. Measuring the similarity of representations

Given a pair of incremental feature extractors \mathcal{F}_0 and \mathcal{F}_N trained by a selected class-incremental algorithm ($N = 5$ for the B500-5step setting and $N = 10$ for the B500-10step setting), we extract a set of features $\{\mathbf{X}_l\}_{l=0}^L$ from \mathcal{F}_0 , where \mathbf{X}_l denotes feature output of layer l . After extracting a corresponding set of features $\{\mathbf{Y}_l\}_{l=0}^L$ from \mathcal{F}_N , we compute CKA¹ between \mathbf{X}_l and \mathbf{Y}_l , $\forall l \in \{0, \dots, L\}$, using the ImageNet-1K \mathcal{D}_0 validation subset. We extract activations from all convolution, batch normalization [11], and residual block layers of ResNet-18, which results in two sets of features, each with a cardinality of 50 ($L = 49$). A

¹More precisely, we use the mini-batch version of CKA, which has been shown to converge to the same value of full-batch CKA, described in Eq. (3) [18]. We provide details for mini-batch CKA in the Appendix.

high CKA between \mathbf{X}_l and \mathbf{Y}_l , *i.e.*, $\text{CKA}(\mathbf{X}_l, \mathbf{Y}_l) \approx 1$, indicates that the two feature representations are highly similar, and thus, remained static across incremental stages. By extension, high CKA may also serve as a strong indicator for *low forgetting* due to high similarity of \mathbf{X}_l and \mathbf{Y}_l .

Figure 3 presents the same-layer CKA between \mathcal{F}_0 and \mathcal{F}_5 trained on various CIL algorithms. Each subplot visualizes the layer-wise CKA for two algorithms: 1) the Naive method, which represents a fully-plastic baseline, and 2) the specified CIL algorithm. We first inspect the naive model, which displays relatively high CKA in early layers, but significantly deteriorates in latter layers. This observation is consistent with the notion that early layers learn low-level and widespread features, while higher layers tend to learn more class-specific information [32]. Then, a cursory examination of all other algorithms suggest that all compared CIL algorithms do indeed enforce feature representations to be similar across incremental models, albeit with varying levels of success. For example, both POD and AANet retain high feature similarities across most layers, although there are some significant drops in a few select layers (more details on CKA of AANet provided in the Appendix). Furthermore, we observe that AFC maintains high CKA across all layers, implying that the model trained with AFC has high stability, and thus, undergoes little to no forgetting in each incremental step.

Finally, we present two plots for DER. In the B500-5step setting, DER consists of 6 separate feature extractors, the first of which is identical to \mathcal{F}_0 . Thus, we plot CKA separately for the fixed feature extractor, DER (0), and all other feature extractors, DER (1~5). As expected, we observe maximum CKA for DER (0), and much lower CKA for DER (1 ~ 5).

4.3. t-SNE Visualizations

To further corroborate our observations from Section 4.2, we visualize the feature shift t-SNE [29]. We randomly sample 5 classes from \mathcal{C}_5 and 20 images from each selected class. Then, we compute the feature representations of each image using both \mathcal{F}_0 and \mathcal{F}_5 , and visualize all 200 features in a single t-SNE plot.

Figure 4 illustrates the t-SNE plots of the Naive and Oracle feature extractors, as well as those of four CIL algorithms that exhibit high stability. In the Naive and Oracle plots, we observe that the same-class features have shifted significantly, since the \mathcal{F}_0 and \mathcal{F}_5 features are clustered in different regions. This is expected since both models exhibit high plasticity. In the AFC, AANet, POD, and LUCIR plots, however, the feature shift is trivial. Interestingly, most of the \mathcal{F}_0 and \mathcal{F}_5 features corresponding to the same inputs are overlapped, even though \mathcal{F}_5 has been trained on $\{\mathcal{D}_i\}_{i=1}^5$, while \mathcal{F}_0 has not. Once again, this suggests that these algorithms excel at maintaining important knowledge

from previous tasks, but fail to acquire new knowledge from incremental data.

5. Improving Class-Incremental Learning

We now outline two methods to improve CIL models, inspired by the analyses presented in Sections 3 and 4.

5.1. Partial-DER

Figures 1 and 2 show that DER not only retains knowledge from old classes, but is also adept at learning concepts from new classes. Such properties of DER are desirable for CIL algorithms. DER achieves this by freezing and maintaining $\{\mathcal{F}_j\}_{j=0}^{i-1}$ at the i^{th} incremental stage, and adding a new fully plastic feature extractor (with the same number of parameters as \mathcal{F}_0) for the current set of classes. This allows the feature extractor to fully learn concepts of new classes all the while maintaining knowledge of old classes. However, a major issue of DER is scalability, since for an N -step CIL model, a single forward pass through \mathcal{F}_N requires $N + 1$ times computation compared to \mathcal{F}_0 .

Based on our CKA analysis in Figure 3, we propose a modification of DER, called partial-DER (pDER), which not only makes DER much more efficient but also improves the overall performance as well. In Section 4.2, we observed that even the Naive model maintains high feature similarity in the lower layers of the network, *i.e.*, the feature representations of lower layers do not change much across \mathcal{F}_0 and \mathcal{F}_5 even when the model is fully plastic. This suggests that lower layers are inherently stable. Thus, instead of maintaining N full feature extractors for an N -step setting, pDER fixes the lower subset of layers in \mathcal{F}_0 , and only applies DER on the upper subset of layers. More specifically, we consider all layers up to ResNet’s Layer 4 as the lower subset, and apply DER only for Layer 4. We find that this simple modification reduces the GMACs of a forward pass through a DER model by up to 65%², while improving $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$ and $\text{Acc}(\mathcal{M}_5, \mathcal{D})$ by 1.5%p and 0.9%p, respectively.

5.2. Exploiting static feature extractors

The second method, which we name “Exploit”, is based on the observation that the feature extractors in most CIL models remain static over the course of incremental stages. In such case, we can significantly improve the training efficiency, achieve strong performance, and eliminate the need for previous-class exemplars by simply freezing the base feature extractor, \mathcal{F}_0 . This exploit can be interpreted as an extreme case of pDER, where branching occurs in the classification layer.

² \mathcal{F}_5 of DER requires 10.9 GMACs for one forward pass while \mathcal{F}_5 of pDER does 3.9 GMACs. More details are provided in the Appendix.

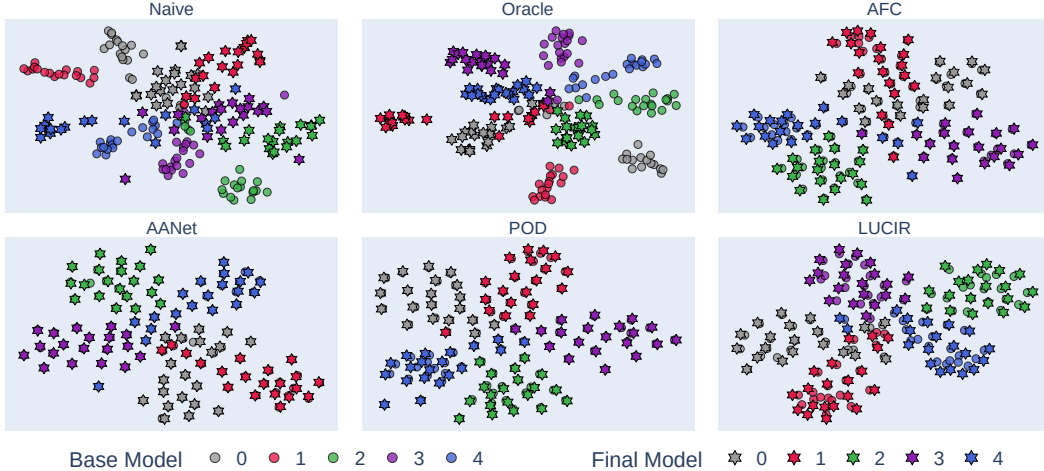


Figure 4. t-SNE visualization of features from \mathcal{F}_0 and \mathcal{F}_5 using 5 classes of \mathcal{C}_5 and 20 samples from each class. All 200 embeddings are visualized in a single plot. Each color represents a distinct class, and features of \mathcal{F}_0 and \mathcal{F}_5 are depicted with a circle and star, respectively.

We first train \mathcal{M}_0 on \mathcal{D}_0 , and fix the feature extractor, \mathcal{F}_0 , for all subsequent incremental steps. The weight matrix of the cosine classifier \mathcal{G} is denoted as $\mathbf{W}_0 \in \mathbb{R}^{F \times |\mathcal{C}_0|}$, where F is the output feature dimension of the feature extractor. For each incremental step i , we train a new weight matrix for the classifier, $\mathbf{W}_i \in \mathbb{R}^{F \times |\mathcal{C}_i|}$, where the cross-entropy loss is only computed for \mathcal{C}_i , *i.e.*, the softmax operation only considers the logits for $|\mathcal{C}_i|$ classes. Since \mathcal{F}_0 is fixed for all stages, we do not need to update $\mathbf{W}_j, \forall j < i$. After training on \mathcal{D}_N , we concatenate the set of weight matrices $\{\mathbf{W}_i\}_{i=0}^N$ to produce a single weight matrix, $\mathbf{W} \in \mathbb{R}^{F \times |\mathcal{C}|}$. Finally, we compose \mathcal{G} , parametrized by weight \mathbf{W} , with \mathcal{F}_0 to obtain the final model: $\mathcal{M}_N = \mathcal{G} \circ \mathcal{F}_0$.

Compared to CIL algorithms that usually train \mathcal{M} for 90 to 120 epochs on each \mathcal{D}_i , our exploit only requires around 10 epochs of training to converge for each incremental stage. Furthermore, we only need to compute gradients for \mathcal{G} , which reduces the computational burden of training. These two factors make training extremely fast compared to traditional algorithms that tune the entire model for many epochs. Despite requiring only a fraction of the computation, our exploit achieves final ImageNet-1K accuracy of 60.5% and an average incremental accuracy of 67.2%, which are competitive against the other methods.

6. Discussions

Based on the plots in Figure 1, we define a metric as

$$\Delta \mathcal{M}'_i := \text{Acc}(\mathcal{M}'_i, \mathcal{D}) - \text{Acc}(\mathcal{M}'_0, \mathcal{D}), \quad (4)$$

which measures the relative performance of improvement of \mathcal{F}_i over \mathcal{F}_0 , *i.e.*, how much the feature extractor improves as training classes are added. A high positive $\Delta \mathcal{M}'_i$ indicates that the feature extractor is able to learn new concepts incrementally. On the other hand, $\Delta \mathcal{M}'_i \simeq 0$ indicates that the

feature extractor is stable, but not plastic enough to acquire new knowledge. Finally, a large negative $\Delta \mathcal{M}'_i$ represents severe catastrophic forgetting in the feature extractor. Ultimately, CIL models should strive to maximize $\Delta \mathcal{M}'_i$ in order to facilitate better feature representation learning.

Are we learning continually? Table 1 quantitatively summarizes all the compared CIL algorithms in terms of $\text{Acc}(\mathcal{M}'_0, \mathcal{D})$, $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$, and $\Delta \mathcal{M}'_5$. Looking at LUCIR, SSIL, POD, AANet, and AFC in Table 1, we notice that there is almost no difference between $\text{Acc}(\mathcal{M}'_0, \mathcal{D})$ and $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$, which all lie in the range of 62% ~ 63%. Furthermore, there exists a significant gap between $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$ of the aforementioned algorithms and $\text{Acc}(\mathcal{M}'_5, \mathcal{D}) = 70.8\%$ of the Oracle model. This implies that, while these methods do mitigate catastrophic forgetting to varying degrees, they suffer from a lack of plasticity, *i.e.*, are unable to learn continually. Moreover, these results also imply that if feature representations do not accumulate new knowledge through incremental data, the highest achievable performance is still 7~8%p shy of the Oracle model even with an optimal classifier. Thus, the plasticity of feature representations is an absolutely crucial aspect of modern incremental learning algorithms.

Effect of task similarity Task similarity has been shown to affect the amount of catastrophic forgetting that neural networks experience [4]. Despite this, we find that our analyses lead to the same conclusions even when the level of task similarity changes across stages. These experimental results can be found in Section 5 of the Appendix.

Pretrained setting At this point it is reasonable to question whether the lack of plasticity may be attributed to the experimental setting, *i.e.* B500-5step, where models are

Table 1. Summary of all compared CIL algorithms. *Methods introduced in Section 5

| Method | Naive | iCARL | LUCIR | SSIL | POD | AAANet | AFC | DER | Exploit* | pDER* | Oracle |
|---|-------|-------|-------|------|------|--------|------|------|----------|-------|--------|
| $\text{Acc}(\mathcal{M}'_0, \mathcal{D})$ | 60.4 | 61.8 | 62.1 | 62.5 | 62.9 | 62.5 | 62.9 | 60.0 | 62.9 | 60.0 | 60.4 |
| $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$ | 52.2 | 53.9 | 62.8 | 63.5 | 62.6 | 62.8 | 62.9 | 66.9 | 62.9 | 68.4 | 70.8 |
| $\Delta\mathcal{M}'_5$ | -8.2 | -7.9 | 0.7 | 1.0 | -0.3 | 0.3 | 0.0 | 6.9 | 0.0 | 8.4 | 10.4 |

Table 2. Average incremental accuracy and Final ImageNet accuracy of all compared CIL algorithms. [†]SSIL uses a different class ordering than all other methods; thus, inter-method comparisons with SSIL may not be appropriate. *Methods introduced in Section 5

| Method | Naive | iCARL | LUCIR | SSIL [†] | POD | AAANet | AFC | DER | Exploit* | pDER* | Oracle |
|--|-------|-------|-------|-------------------|------|--------|------|------|----------|-------|--------|
| Avg. Inc. Acc. | 41.4 | 31.5 | 55.6 | 65.8 | 66.3 | 57.1 | 66.4 | 69.1 | 67.2 | 69.7 | 72.8 |
| $\text{Acc}(\mathcal{M}_5, \mathcal{D})$ | 31.5 | 17.2 | 41.0 | 59.8 | 58.6 | 43.3 | 59.0 | 63.8 | 60.5 | 64.7 | 70.8 |

pretrained with 500 classes before being updated incrementally. While this may be true to some extent³, we argue that continual learners should strive towards obtaining better feature representations *even if* the pretrained representations are somewhat sufficient. If not, what good is a continual model compared to a static one as described in Section 5.2? In fact, we have shown that the pretrained feature extractor is still lackluster compared to the Oracle while DER has shown that, albeit lacking scalability, it is possible to continually learn stronger representations even with a pretrained feature extractor. Thus, our choice to conduct analysis on this pretrained setting cannot justify the lack of plasticity exhibited by most CIL algorithms.

Limitations of traditional metrics In Table 2 we present the widely adopted metrics in CIL: 1) average incremental accuracy of all compared CIL algorithms on the B500-5step setting and 2) the final accuracy on ImageNet, $\text{Acc}(\mathcal{M}_5, \mathcal{D})$. In particular, we note that our Exploit in Table 2 outperforms all but one CIL algorithm (DER) on both metrics, despite keeping the feature extractor \mathcal{F}_0 fixed. However, according to Table 1, all the aforementioned CIL algorithms demonstrate similar $\text{Acc}(\mathcal{M}'_0, \mathcal{D})$ and $\text{Acc}(\mathcal{M}'_5, \mathcal{D})$ scores. This suggests that while the feature representations of all compared methods except Naive, iCARL, and DER have similar levels of discriminativeness, it is not well expressed in terms of average incremental accuracy nor $\text{Acc}(\mathcal{M}_5, \mathcal{D})$. Clearly, high average incremental accuracy and final accuracy are not really indicative of how much the model has *learned* continually; yet, these metrics have become the de-facto standard in CIL research. This should be alarming for both researchers and practitioners; ambiguous metrics deliver a false illusion of progress, and may lead researchers to develop algorithms that *seem* to outperform other state-of-the-art algorithms, but are completely misaligned with the motivation behind continual learning, *e.g.* our Exploit. Thus, we hope that the analyses in our work will facilitate better evaluation of CIL al-

³Analysis of select algorithms on B0-10step included in the Appendix.

gorithms and inspire researchers to focus more on stronger feature representation learning of incremental models.

Connection with theory Of the works we study, the only algorithm that exhibits both high plasticity and stability is based on parameter isolation, which makes us wonder whether it is even possible to achieve a strong balance of stability and plasticity without adding new model parameters. However, Raghavan *et al.* [24] prove that there can exist a stable equilibrium point between forgetting and generalization. In fact, it is possible that the low-plasticity methods do indeed reach a stable equilibrium, except that this equilibrium is less optimal in terms of performance. Therefore, a key question for future works is: how can we reach a *more optimal* equilibrium point at each incremental stage without adding more parameters?

7. Conclusion

We took a deep dive into how effectively modern CIL algorithms address the stability-plasticity dilemma. We introduced evaluation protocols that help us better understand the stability and plasticity of feature representations. Our evaluations of recent works showed that many CIL methods are too fixated on the notion of alleviating catastrophic forgetting, to the extent that the feature extractor rarely learns any new concepts after the initial stage of training, namely on \mathcal{D}_0 . Based on this observation, we introduced two simple algorithms that improve upon an existing algorithm and exploit the shortcomings of the standard evaluation metrics for CIL research. All in all, we hope that our findings will propel CIL research to focus more on stronger continual learning of feature representations.

Acknowledgements This work was partly supported by Samsung Electronics Co., Ltd., and by the NRF Korea grant [No. 2022R1A2C3012210] and the IITP grants [No.2022-0-00959; 2021-0-02068; 2021-0-01343] funded by the Korean government (MSIT).

References

- [1] Hongjoon Ahn, Jihwan Kwak, Subin Lim, Hyeonsu Bang, Hyojun Kim, and Taesup Moon. SS-IL: separated softmax for incremental learning. In *ICCV*, 2021. 1, 2, 3
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Mach. Learn.*, 2010. 1
- [3] Corinna Cortes, Mehryar Mohri, and Afshin Rostamizadeh. Algorithms for learning kernels based on centered alignment. *JMLR*, 2012. 1, 5
- [4] T. Doan et al. A theoretical analysis of catastrophic forgetting through the ntk overlap matrix. In *AISTATS*, 2021. 7
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 5
- [6] Arthur Douillard, Matthieu Cord, Charles Ollion, and Thomas Robert. PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning. In *ECCV*, 2020. 1, 3
- [7] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015. 1
- [8] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. In *Algorithmic Learning Theory*, 2005. 5
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 2
- [10] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a Unified Classifier Incrementally via Rebalancing. In *CVPR*, 2019. 1, 2, 3
- [11] Sergey Ioffe and Christian Szegedy. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 2015. 5
- [12] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *CVPR*, 2022. 1, 3
- [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*, 2017. 1
- [14] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *ICML*, 2019. 1, 5
- [15] Seungmin Lee, Dongwan Kim, Namil Kim, and Seong-Gyun Jeong. Drop to adapt: Learning discriminative features for unsupervised domain adaptation. In *ICCV*, 2019. 1
- [16] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive Aggregation Networks for Class-Incremental Learning. In *CVPR*, 2021. 1, 2
- [17] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *CVPR*, 2016. 1
- [18] Thao Nguyen, Maithra Raghu, and Simon Kornblith. Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. In *ICLR*, 2021. 5
- [19] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning Deconvolution Network for Semantic Segmentation. In *ICLR*, 2015. 1
- [20] Dongmin Park, Seokil Hong, Bohyung Han, and Kyoung Mu Lee. Continual learning by asymmetric loss approximation with single-side overestimation. In *ICCV*, 2019. 1
- [21] Jaeyoo Park, Minsoo Kang, and Bohyung Han. Class-incremental learning for action recognition in videos. In *ICCV*, 2021. 1
- [22] Ameya Prabhu, Philip HS Torr, and Puneet K Dokania. GDumb: A Simple Approach that Questions Our Progress in Continual Learning. In *ECCV*, 2020. 2
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1
- [24] K. Raghavan and P. Balaprakash. Formalizing the generalization-forgetting trade-off in continual learning. In *NeurIPS*, 2021. 8
- [25] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, 2021. 5
- [26] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental Classifier and Representation Learning. In *CVPR*, 2017. 1, 2
- [27] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. In *CVPR*, 2016. 1
- [28] 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. 2
- [29] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 2008. 1, 6
- [30] 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
- [31] Shipeng Yan, Jiangwei Xie, and Xuming He. DER: Dynamically Expandable Representation for Class Incremental Learning. In *CVPR*, 2021. 1, 3
- [32] Matthew D. Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014. 6
- [33] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shutao Xia. Maintaining Discrimination and Fairness in Class incremental Learning. In *CVPR*, 2020. 3