

Always Be Dreaming: A New Approach for Data-Free Class-Incremental Learning

James Smith^{1*}, Yen-Chang Hsu², Jonathan Balloch¹, Yilin Shen², Hongxia Jin², Zsolt Kira¹

¹Georgia Institute of Technology, ²Samsung Research America

Abstract

Modern computer vision applications suffer from catastrophic forgetting when incrementally learning new concepts over time. The most successful approaches to alleviate this forgetting require extensive replay of previously seen data, which is problematic when memory constraints or data legality concerns exist. In this work, we consider the high-impact problem of Data-Free Class-Incremental Learning (DFCIL), where an incremental learning agent must learn new concepts over time without storing generators or training data from past tasks. One approach for DFCIL is to replay synthetic images produced by inverting a frozen copy of the learner’s classification model, but we show this approach fails for common class-incremental benchmarks when using standard distillation strategies. We diagnose the cause of this failure and propose a novel incremental distillation strategy for DFCIL, contributing a modified cross-entropy training and importance-weighted feature distillation, and show that our method results in up to a 25.1% increase in final task accuracy (absolute difference) compared to SOTA DFCIL methods for common class-incremental benchmarks. Our method even outperforms several standard replay based methods which store a coreset of images. Our code is available at <https://github.com/GT-RIPL/AlwaysBeDreaming-DFCIL>

1. Introduction

A shortcoming of modern computer vision settings is that they often assume offline training with a large dataset encompassing all objects to be encountered during deployment. In practice, many applications require a model be continuously updated after new environments/situations are encountered. This is the *class-incremental learning* paradigm (also known as a subset of continual or lifelong learning), with the loss of knowledge over sequences of learning tasks referred to as *catastrophic forgetting*. Suc-

cessful incremental learning approaches have an unfortunate commonality: they require extensive memory for replay of previously seen or modeled data to avoid the catastrophic forgetting problem. This is concerning for many computer vision applications because 1) Many computer vision applications are on-device and therefore memory constrained [15, 35, 52], and 2) Many computer vision applications learn from data which cannot be legally stored [5, 10, 62]. This leads us to ask: *How can computer vision systems incrementally incorporate new information without storing data?* We refer to this setting as *Data-Free Class-Incremental Learning (DFCIL)* (also known as Data-Free Continual Learning [60]).

An intuitive approach for DFCIL is to simultaneously train a generative model to be sampled for replay [25, 26, 49, 58]. Unfortunately, training a generative model is much more computationally and memory intensive compared to a classification model. Additionally, it is not clear whether generating images from the data distribution will violate data legality concerns because using a generative model increases the chance of memorizing potentially sensitive data [42]. Instead, we explore the concept of *model-inversion image synthesis*, where we can invert the already provided inference network to obtain images with similar activations in the network to the training data. This idea is inviting because it requires no additional networks to be trained (it only requires the existing inference network) and is less susceptible to data privacy concerns.

The closest existing work for the DFCIL problem is DeepInversion [60], which optimizes random noise into images for knowledge distillation using a frozen teacher network. DeepInversion is designed for standard student-teacher knowledge distillation and achieves state-of-the-art performance for this task. Unfortunately, the authors report that when trying class-incremental learning for tasks where the old images and new images are similar (such as tasks from the same dataset, a standard benchmarking practice for class-incremental learning), their method performs “statistically equivalent or slightly worse compared to Learning without Forgetting (LwF)”, with LwF [34] being their most competitive existing baseline.

*Correspondence to: James Smith jamessealesmith@gatech.edu

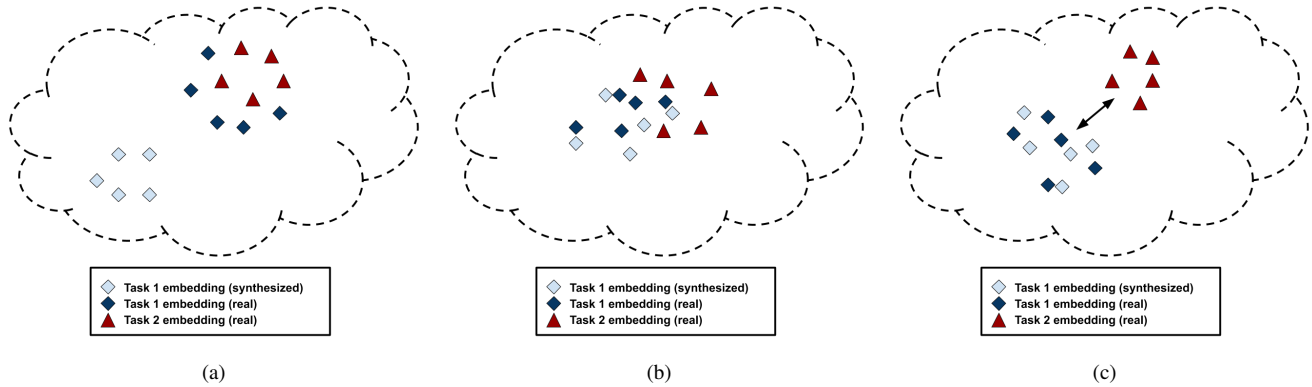


Figure 1: The distribution of feature embeddings when using synthetic replay data for class-incremental learning. (a) A straight application of synthetic data makes the model learn features more distinguishable between real and fake instead of task 1 and 2. This is the main problem analyzed and addressed in this work. (b) Modifying classification loss and adding regularization mitigates the feature drifting between real and fake. (c) This is the desired feature distributions. Our method makes task 1 and 2 more separable.

The goal of this paper (summarized in Figure 1) is to dissect the cause of this failure and propose a solution for DFCIL. Specifically, we reason that when training a model with real images from the current task and synthetic images representing the past tasks, the feature extraction model causes the feature distributions of **real images from the past tasks** (which are not available during training) to be *close* in the feature space to the **real images from the current task** and *far* in the feature space from the **synthetic images**. This causes a bias for the model to falsely predict real images from the previous tasks with current task labels. This phenomena indicates that when training a network with two distributions of data, containing both a semantic shift (past tasks versus current task) and a distribution shift (synthetic data versus real data), the distribution shift has a higher effect on the feature embeddings. Thus, validation/test images from the previous classes will be identified as new classes due to the model fixating on their domain (i.e., realistic versus synthetic pixel distribution) rather than their semantic content (i.e., past versus current task).

To address this issue, we propose a novel class-incremental learning method which learns features for the new task with a *local* classification loss which excludes the synthetic data and past-task linear heads, instead relying on importance-weighted feature distillation and linear head fine-tuning to separate feature embeddings of the new and past tasks. We show that our method represents the new state of the art for the DFCIL setting, resulting in up to a 25.1% increase in final task accuracy (absolute difference) compared to DeepInversion for common class-incremental benchmarks, and even outperforms popular replay baselines Naive Rehearsal and LwF with a coreset. *In summary, we make the following contributions:*

1. We use a classic class-incremental learning benchmark to diagnose and analyze why standard distillation approaches for class-incremental learning (such as Deep-

Inversion) fail when using synthetic replay data.

2. We directly address this failure with a modified cross-entropy minimization, importance-weighted feature distillation, and linear head fine-tuning.
3. We achieve a new state of the art performance for the DFCIL setting.

2. Background and Related Work

Catastrophic Forgetting: Approaches to mitigate catastrophic forgetting can be organized into a few broad categories and are all useful depending on which constraints are present. For example, methods which expand a model’s architecture as new tasks are encountered are highly effective for applications where a model growing with tasks is practical [12, 32, 36, 40, 48, 51]. We do not consider these methods because the model parameters grow linearly with the number of tasks. Experience replay with stored data [2, 3, 7, 8, 13, 17, 22, 27, 38, 44, 45, 46, 57] or samples from a generative model [25, 26, 43, 49, 54] is highly effective when storing training data or training/saving a generative model is possible, unlike the DFCIL setting.

Another approach is to regularize the model with respect to past task knowledge while training the new task. This can either be done by regularizing the model in the weight space (i.e., penalize changes to model parameters) [1, 11, 28, 53, 61] or the prediction space (i.e., penalize changes to model predictions) [6, 21, 31, 34, 50]. Prediction space regularization (accomplished using *knowledge distillation*) has been found to perform better than model regularization based methods for class-incremental learning [33, 55].

Knowledge Distillation in Class-Incremental Learning: Based on the original work proposing knowledge distillation from a large model to a smaller model [20], methods such as Learning Without Forgetting [34], Distillation and Retrospection [21], End-To-End Incremental Learning [6],

Global Distillation [31], and Bias Correction [59] have effectively leveraged knowledge distillation as a prediction regularization technique for incremental learning. The high level idea of knowledge distillation is to periodically save a frozen copy of the model (here we use a ConvNet) and to ensure the new model makes similar predictions to the frozen model over a set of distillation images (while simultaneously learning the new task). Knowledge distillation does not require that the frozen model be replaced at the task sequence boundaries, but this is typically done when evaluating competing methods. This regularization can also take place in the feature space rather than the prediction space [4, 22], which we refer to as *feature distillation*. These knowledge distillation methods require stored data over which to enforce similar predictions, but the next section describes a form of knowledge distillation which does not require training data.

Data-Free Knowledge Distillation: Knowledge from a neural network can be transferred in the absence of training data. We refer to the line of work which *synthesizes* distillation images using the trained inference network itself and resulting activation statistics as *data-free knowledge distillation*. This approach is very important for applications where training data is sensitive and not easily available for legality issues. The first such work we are aware of, DeepDream [41], optimizes randomly generated noise into images which minimize classification loss and an image prior. Another early method [37] matches stored layer statistics from a trained “teacher” model while leaving a small memory footprint using frequency-based compression techniques. The Data-Free Learning method [9] exploits a GAN architecture to synthesize images which match the trained teacher’s statistics while balancing content losses which maximize both temperature scaled linear heads (to drive class-specific content high) and class prediction entropy (to encourage high diversity of classes sampled). Three recent methods leverage layer content stored in batch-normalization layers to synthesize realistic looking images for knowledge distillation [16, 39, 60].

To the best of our knowledge, only two class-incremental learning methods are designed for the data-free perspective. Automatic-Recall Machines (ARM) [24] perturb training data from the current task into images which maximize “forgetting” from the past tasks. However, this method is designed for a “single pass” setting where data is only trained for one epoch, which is a different setting than ours. DeepInversion [60] also evaluates data-free knowledge distillation in the class-incremental learning paradigm, but only found success in small task sequences (max of three) using tasks which are very distinct in image content. Our paper dissects why the DeepInversion method fails at difficult class-incremental learning problems and proposes a solution for successful data-free class-incremental learning.

3. Preliminaries

Class-Incremental Learning: In class-incremental learning, a model is shown labeled data corresponding to M semantic object classes c_1, c_2, \dots, c_M over a series of N tasks corresponding to non-overlapping subsets of classes. We use the notation \mathcal{T}_n to denote the set of classes introduced in task n , with $|\mathcal{T}_n|$ denoting the number of object classes in task n . Each class appears in only a single task, and the goal is to incrementally learn to classify new object classes as they are introduced while retaining performance on previously learned classes. The class-incremental learning setting is challenging because no task indexes are provided to the learner during inference and the learner must support classification across all classes seen up to task n [23]. This is more difficult than *task-incremental learning*, where the task indexes are given during both training and inference. While our setting does not necessitate known task boundaries during training, we follow prior works [34, 23, 60] for fair comparison and create model copies at the task boundaries for each method.

To describe our inference model, we denote $\theta_{i,n}$ as the model θ at time i that has been trained with the classes from task n . For example, $\theta_{n,1:n}$ refers to the model trained during task n and its logits associated with all tasks up to and including class n . We drop the second index when describing the model trained during task n with all logits (for example, θ_n).

4. Baseline Approach for Data-Free Class Incremental Learning

In this section, we propose a general baseline for data-free class incremental learning based on efforts in prior work. We start by summarizing the *data synthesis* (i.e. generation of images from the discriminative model itself) approach we found most successful for class-incremental learning. We then review pertinent *knowledge distillation* losses, ultimately focusing on the loss functions used by DeepInversion [60] for class-incremental learning.

Model-Inversion Image Synthesis: Most model-inversion image synthesis approaches seek to synthesize images by directly optimizing them with respect to a prior discriminative model θ_{n-1} . For B synthetic images, a tensor $\hat{X} \in \mathcal{R}^{B \times H \times W \times C} = \{\hat{x}_1 \dots \hat{x}_b\}$, where H , W , and C correspond to the training data image dimensions, is initialized from Gaussian noise. However, it is computationally inefficient to optimize one batch of images at a time. Especially considering that class-incremental learning is expected to be computationally efficient, we choose to approximate this optimization from noise to synthesized images using a ConvNet-parameterized function \mathcal{F}_ϕ . This allows the framework to train \mathcal{F}_ϕ once per task (using *only* θ_{n-1} , i.e., no data), only store it temporarily during that given

task, sample synthetic images as needed during training for task \mathcal{T}_n , and then discarded it at the end of the task.

\mathcal{F}_ϕ can struggle with synthetic class diversity; rather than condition \mathcal{F}_ϕ on class labels Y , we follow [9] and optimize the diversity of class predictions of synthetic images to match a uniform distribution. Denoting $p_\theta(x)$ as the predicted class distribution produced by model θ for some input x , we want to *maximize* the entropy of the mean class prediction vector for synthetic samples \hat{X} . Formally, we minimize a *label diversity loss*:

$$\mathcal{L}_{div}(\hat{Y}) = -\mathcal{H}_{info} \left(\frac{1}{B} \sum_b p_\theta(\hat{x}_b) \right) \quad (1)$$

where \mathcal{H}_{info} is the information entropy. Notice when the loss is taken at the minimum, every element in the mean class prediction vector would equal $\frac{1}{|\mathcal{T}_1 \dots \mathcal{T}_{n-1}|}$, meaning that classes are generated with at roughly the same rate.

In addition to diversity, to consistently synthesize useful images in the DFCIL setting, the images must enforce calibrated class confidences, consistency of feature statistics, and a locally smooth latent space, described below.

Content loss, \mathcal{L}_{con} , maximizes class prediction confidence with respect to the image tensor \hat{X} such that θ_{n-1} should make confident predictions on all inputs. Formally, \mathcal{L}_{con} is the cross entropy classification loss between the class predictions of \hat{X} and the maximum class predictions \hat{Y} :

$$\mathcal{L}_{con}(\hat{X}, \hat{Y}) = \mathcal{L}_{CE} \left(p_{\theta_{n-1,1:n-1}}^{\alpha_{temp}}(\hat{x}), \hat{y} \right) \quad (2)$$

$$\hat{y} = \underset{\hat{y} \in \mathcal{T}_1 \dots \mathcal{T}_{n-1}}{\operatorname{argmax}} p_{\theta_{n-1,1:n-1}}(\hat{x}) \quad (3)$$

where \mathcal{L}_{CE} is standard cross entropy loss and the logit output of θ is scaled by a temperature constant α_{temp} . By combining \mathcal{L}_{con} with \mathcal{L}_{div} , we ensure that the synthesized images will represent the distribution of all past task classes

Prior work has found that the complexity of model-inversion can cause the distribution of θ_{n-1} features to deviate greatly from distributions over batches of synthetic images. Intuitively, the batch statistics of synthesized images should match those of the batch normalization layers in θ_{n-1} . To enforce this, *stat alignment loss*, \mathcal{L}_{stat} , penalizes the deviation between intermediate layer batch-normalization statistics (BNS) stored in θ_{n-1} and features at those layers for synthetic images [16, 39, 60]:

$$\mathcal{L}_{stat}(\hat{X}) = \frac{1}{L} \sum_{l=1}^L BNS(\hat{\mu}_{\hat{X},l}, \hat{\sigma}_{\hat{X},l}, \mu_l, \sigma_l) \quad (4)$$

$$\begin{aligned} BNS(\hat{\mu}, \hat{\sigma}, \mu, \sigma) &= KL(\mathcal{N}(\mu, \sigma^2) || \mathcal{N}(\hat{\mu}, \hat{\sigma}^2)) \\ &= \log \frac{\hat{\sigma}}{\sigma} - \frac{1}{2} \left(1 - \frac{\sigma^2 + (\mu - \hat{\mu})^2}{\hat{\sigma}^2} \right) \end{aligned} \quad (5)$$

where KL denotes the Kullback-Leibler (KL) divergence, $\hat{\mu}_{\hat{X},l}, \hat{\sigma}_{\hat{X},l}$ are the mean and standard deviation of the features at layer l for a given mini-batch of synthesized image, and μ_l, σ_l are the batch-norm statistics of said layer l . Because batch statistics of θ_{n-1} are stored in batch normalization layers, this loss does not require any additional storage.

Additionally, prior knowledge tells us natural images are more locally smooth in pixel space than the initial noise. As such, we can stabilize the optimization by minimizing the *smoothness prior loss* \mathcal{L}_{prior} . Formally, \mathcal{L}_{prior} is the L2 distance between each synthesized image (\hat{x}) and a version blurred with Gaussian kernel (\hat{x}_{blur}):

$$\mathcal{L}_{prior}(\hat{X}) = \|\hat{x} - \hat{x}_{blur}\|_2^2 \quad (6)$$

All together, assuming the use of \mathcal{F}_ϕ for efficiency, the final loss for the baseline is therefore:

$$\begin{aligned} \min_{\mathcal{F}_\phi} \quad & \alpha_{con} \mathcal{L}_{con}(\hat{X}, \hat{Y}) + \alpha_{div} \mathcal{L}_{div}(\hat{Y}) \\ & + \alpha_{stat} \mathcal{L}_{stat}(\hat{X}) + \alpha_{prior} \mathcal{L}_{prior}(\hat{X}) \end{aligned} \quad (7)$$

Importantly: although we optimize \mathcal{F}_ϕ rather than \hat{X} as done in [16, 60], this method can use the latter with a sacrifice to computational efficiency.

Distilling Synthetic Data for Class-Incremental Learning:

In the class-incremental learning setting, where the classes of task \mathcal{T}_n are modeled without unlearning the representation of classes of tasks $\mathcal{T}_1 \dots \mathcal{T}_{n-1}$, knowledge distillation over the synthesized images is most often used to regularizes θ_n , forcing it to learn \mathcal{T}_n with minimal degradation to $\mathcal{T}_1 \dots \mathcal{T}_{n-1}$ knowledge. For task \mathcal{T}_n , we synthesize images from a frozen copy of (θ_{n-1}) trained during task \mathcal{T}_{n-1} . These synthetic images then help us distill knowledge learned in tasks $\mathcal{T}_1 \dots \mathcal{T}_{n-1}$ into our current model (θ_n) as it learns from \mathcal{T}_n data.

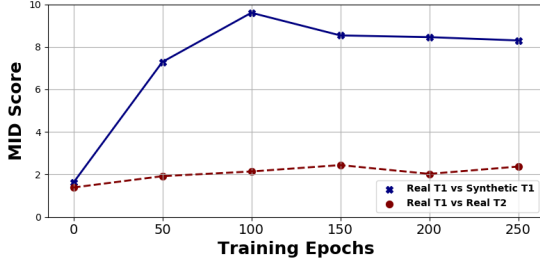
In our baseline, we adopt the distillation approach used in DeepInversion [60], which generalizes the original Learning without Forgetting (LwF) [34] distillation approach. Formally, given current task data, x , and synthesized distillation data, \hat{x} , we minimize:

$$\begin{aligned} \min_{\theta_n} \quad & \mathcal{L}_{CE}(p_{\theta_{n,1:n}}(x), y) + \mathcal{L}_{KD}^{DI}(x, \theta_n, \theta_{n-1}) \\ & + \mathcal{L}_{KD}^{DI}(\hat{x}, \theta_n, \theta_{n-1}) \end{aligned} \quad (8)$$

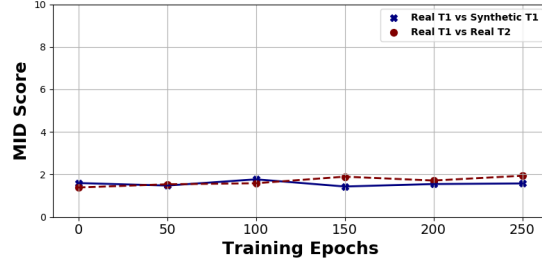
where \mathcal{L}_{KD}^{DI} is a knowledge distillation regularization as:

$$\mathcal{L}_{KD}^{DI}(x, \theta_n, \theta_{n-1}) = KL(p_{\theta_{n-1,1:n}}(x) || p_{\theta_{n,1:n}}(x)) \quad (9)$$

Here, $p_{\theta_{n-1,1:n}}(x)$ is simply $p_{\theta_{n-1,1:n-1}}(x)$ appended with zeros to represent zero class-probability for the new classes (which are not available for $\theta_{n-1,1:n-1}$). The key idea of \mathcal{L}_{KD}^{DI} is that, as the logits of the teacher and student will always be different in the class incremental learning setting, appending zeros to the class-probability vectors aligns the student and teacher logit dimensions for a better transfer of knowledge.



(a) DeepInversion [60]



(b) Our Method

Figure 2: Representational distance scores (MID) between feature embeddings of real task 1 data and synthetic task 1 data (blue), real task 2 data (red). Task 1 corresponds to ten classes of CIFAR-100 while task 2 corresponds to a different ten classes of CIFAR-100; the results are generated after training on task 2.

5. Diagnosis: Feature Embedding Prioritizes Domains Over Semantics

To understand why the baseline approach for DFCIL fails, we analyze *representational distance* between embedded features with a metric that captures the distance between mean embedded images of two distribution samples. Specifically, we assign a Mean Image Distance (MID) score between a reference sample of images x_a and another sample of images x_b , where a *high score* indicates *dissimilar* features and a *low score* indicates *similar* features. We calculate this score as:

$$MID(z_a, z_b) = \left\| \frac{\mu_a - \mu_b}{\sigma_a} \right\|_2 \quad (10)$$

where z_a, z_b is the penultimate feature embedding of x_a, x_b ; μ_a, μ_b are the mean image feature embeddings of x_a, x_b ; and σ_a^2 is the feature variance of x_a . We normalize the distance between mean embedded images by the standard deviation of the reference distribution sample x_a to minimize the impact of highly deviating features. Additional analysis using Maximum Mean Discrepancy (MMD) [14] is available in our Appendix.

For our analysis, we start by training our model for the first two tasks in the ten-task CIFAR-100 benchmark described in Section 7. We calculate MID between feature embeddings of real task 1 data and real task 2 data, and then we calculate MID between feature embeddings of real task 1 data and synthetic task 1 data. The results are reported in Figure 2. For (a) DeepInversion, the MID score between real task 1 data and synthetic task 1 data is significantly higher than the MID score between real task 1 data and real task 2 data. This indicates that the embedding space prioritizes *domain* over *semantics*, which is detrimental because the classifier will learn the decision boundary between synthetic task 1 and real task 2, introducing great classification error with real task 1 images. This diagnosis motivates our approach, which is proposed in the following section. For (b) our method, the MID score between real task 1 data and

synthetic task 1 data is much lower, indicating that our feature embedding prioritizes *semantics* over *domain*.

6. A New Distillation Strategy for DFCIL

We take the perspective that continual learning should balance: (i) learning features for the new task, (ii) minimizing feature drift over the previous task, and (iii) separating class overlap between new and previous classes in the embedding space (this is similarly discussed in another work [22] under a different setting). Generally, (i) and (iii) are simultaneously achieved with \mathcal{L}_{CE} , but we argue that by separating this into two different losses, features for the new task are learned which do not discriminate between real and synthetic images (i.e. avoid the feature domain bias problem). Following this idea, we propose a new class-incremental learning approach designed for DFCIL which addresses each of these goals independently, as described in the rest of this section.

Learning current task features: The intuition behind our method is to learn features for our current task while circumventing the feature embedding for real data becoming highly biased towards the most recent task. That is, we form \mathcal{L}_{CE} so that the likelihood of x being real versus synthetic, $p_{\theta_{n,n}}(x \in X_{real})$, is not helpful in its minimization. We do this by computing cross entropy classification loss *locally* across the new class linear heads without including the past class linear heads. With this formation, we prevent the model from learning to separate the new and past class data via domain (i.e. synthetic vs. real). Formally, we minimize:

$$\mathcal{L}_{CE}^{\mathcal{T}_n}(p_{\theta_{n,n}}(x), y) = \mathcal{L}_{CE}(p_{\theta_{n,n}}(x|y \in \mathcal{T}_n), y) \quad (11)$$

Minimizing feature drift over previous task data: Because our distillation images are of another domain than the real current task images (causing the feature domain bias problem), we seek a loss function which directly alleviates forgetting in the feature space. An alternative to standard knowledge distillation over softmax predictions (\mathcal{L}_{KD}^{DI}) is *feature distillation*, which instead distills the feature content




Figure 3: Our approach combines (i) learning features for the new task with Eq. (11), (ii) minimizing feature drift over the previous task with Eq. (13), and (iii) separating class overlap between new and previous classes in the embedding space with Eq. (14). We use blue arrows to designate the compute path of the synthetic previous tasks data, green arrows to designate the compute path of the real current task data, and yellow arrows to designated the compute path of both real and synthetic data. We separate out the task \mathcal{T}_n head to show that the local CE loss Eq. (11) uses only this head.

from the penultimate layer. This is formally given as:

$$\mathcal{L}_{KD}^{feat}(\cdot) = \|\theta_{n,1:n-1}^{L-1}(x) - \theta_{n-1,1:n-1}^{L-1}(x)\|_2^2 \quad (12)$$

where $L - 1$ denotes the penultimate layer output of the model. Our intuition is that there exists a trade-off between standard knowledge distillation and feature distillation. \mathcal{L}_{KD}^{feat} reinforces important components of past task data, but it is a strong regularization which inhibits the plasticity of the model (and its ability to learn the new task). On the other hand, \mathcal{L}_{KD}^{DI} does not inhibit learning the new task but can be minimized with a solution in which feature drift has occurred, resulting in the real vs synthetic bias.

Instead, we desire an *importance-weighted* feature distillation which reinforces only the most important components of past task data while allowing less important features to be adapted for the new task. We simply use the linear heads of $\mathcal{T}_1 \cdots \mathcal{T}_{n-1}$ from the frozen model θ_{n-1} , or:

$$\mathcal{L}_{KD}^{wfeat}(\cdot) = \|\mathcal{W}(\theta_{n,1:n-1}^{L-1}(x)) - \mathcal{W}(\theta_{n-1,1:n-1}^{L-1}(x))\|_2^2$$

where $\mathcal{W} = \theta_{n-1,1:n-1}^L$ (13)

By using this importance-weight matrix, features associated with a high magnitude in \mathcal{W} are more important to preserve. In this way, the frozen linear heads from the past tasks indicate approximately how much a change to each feature affects the class distribution.

Separating Current and Past Decision Boundaries: Finally, we need to separate the decision boundaries of the current and past classes without allowing the feature space to distinguish between the real and synthetic data. We do this by fine-tuning the linear classification head of $\theta_{n,1:n}$ with standard cross-entropy loss. Importantly, this loss does not update any parameters in $\theta_{n,1:n}$ besides the final linear classification head. Formally, we minimize:

$$\mathcal{L}_{FT}^{\mathcal{T}_{1:n}}(p_{\theta_{n,1:n}}(x, y)) = \mathcal{L}_{CE}(p_{\theta_{n,1:n}}^*(x, y)) \quad (14)$$

where p^* is calculated with $\theta_{n,1:n}^{1:L-1}$ (i.e., every layer in the model except the classification layer) frozen, updating only $\theta_{n,1:n}^L$. As done in [31, 49, 59], we add task-balancing loss weighting to balance the contributions from the current task with the past tasks.

Final Objective: Visualized in Figure 3, our final optimization objective is given as:

$$\begin{aligned} \min_{\theta_n} \quad & \mathcal{L}_{CE}^{\mathcal{T}_n}(p_{\theta_{n,n}}(x, y)) + \lambda_{kd} \mathcal{L}_{KD}^{wfeat}(\{x, \hat{x}\}, \theta_n, \theta_{n-1}) \\ & + \lambda_{ft} \mathcal{L}_{FT}^{\mathcal{T}_{1:n}}(p_{\theta_{n,1:n}}(\{x, \hat{x}\}, \{y, \hat{y}\})) \end{aligned} \quad (15)$$

where the λ terms weight the contributions of \mathcal{L}_{KD}^{wfeat} and \mathcal{L}_{FT} with respect to $\mathcal{L}_{CE}^{\mathcal{T}_n}$

7. Experiments

We evaluate our approach with several image datasets in the class incremental continual learning setting. We implemented baselines which do not store training data for rehearsal: Deep Generative Replay (DGR) [49], Learning without Forgetting (LwF) [34], and Deep Inversion (Deep-Inversion) [60]. Additionally, we report the upper bound performance (i.e., trained offline) and performance for a neural network trained only on classification loss using the new task training data (we refer to this as Base). We note that a downside of any generative method is that they (1) require long-term storage of a generative model and (2) may violate data legality concerns.

For a fair comparison, our implementation of DeepInversion uses the same image synthesis strategy as our method, with the difference being the distillation method. We do not tune hyperparameters on the full task set because tuning hyperparameters with hold out data from all tasks may violate the principal of continual learning that states each task is visited only once [56]. Importantly, we shuffle the class

Table 1: Results (%) for *data-free* class-incremental learning on CIFAR-100 for various numbers of tasks (5, 10, 20). Results are reported as an average of 3 runs.

Tasks		5		10		20	
Method	Replay Data	A_N (\uparrow)	Ω (\uparrow)	A_N (\uparrow)	Ω (\uparrow)	A_N (\uparrow)	Ω (\uparrow)
Upper Bound	None	69.9 \pm 0.2	100.0 \pm 0.0	69.9 \pm 0.2	100.0 \pm 0.0	69.9 \pm 0.2	100.0 \pm 0.0
Base	None	16.4 \pm 0.4	48.9 \pm 1.1	8.8 \pm 0.1	32.1 \pm 1.1	4.4 \pm 0.3	19.7 \pm 0.7
DGR [49]	Generator	14.4 \pm 0.4	45.5 \pm 0.9	8.1 \pm 0.1	30.5 \pm 0.6	4.1 \pm 0.3	19.0 \pm 0.3
LwF [34]	None	17.0 \pm 0.1	49.5 \pm 0.1	9.2 \pm 0.0	33.3 \pm 0.9	4.7 \pm 0.1	20.1 \pm 0.3
LwF [34]	Synthetic	16.7 \pm 0.1	49.8 \pm 0.1	8.9 \pm 0.0	32.3 \pm 0.0	4.7 \pm 0.0	19.7 \pm 0.0
DeepInversion [60]	Synthetic	18.8 \pm 0.3	53.2 \pm 0.9	10.9 \pm 0.6	37.9 \pm 0.8	5.7 \pm 0.3	23.6 \pm 0.7
Ours	Synthetic	43.9 \pm 0.9	78.6 \pm 1.1	33.7 \pm 1.2	69.6 \pm 1.6	20.0 \pm 1.4	52.5 \pm 2.5

Table 2: Results (%) for class-incremental learning *with replay data* on CIFAR-100 for various numbers of tasks (5, 10, 20). A coreset of 2000 images is leveraged for replay-based methods, and thus *these methods do not meet problem the DFCIL constraints* (note we report for our method numbers *without* any coreset). Results are reported as an average of 3 runs.

Tasks		5		10		20	
Method	Replay Data	A_N (\uparrow)	Ω (\uparrow)	A_N (\uparrow)	Ω (\uparrow)	A_N (\uparrow)	Ω (\uparrow)
Upper Bound	None	69.9 \pm 0.2	100.0 \pm 0.0	69.9 \pm 0.2	100.0 \pm 0.0	69.9 \pm 0.2	100.0 \pm 0.0
Naive Rehearsal	Coreset	34.0 \pm 0.2	73.4 \pm 0.8	24.0 \pm 1.0	64.6 \pm 2.1	14.9 \pm 0.7	51.4 \pm 2.9
LwF [34]	Coreset	39.4 \pm 0.3	79.0 \pm 0.0	27.4 \pm 0.8	69.4 \pm 0.4	16.6 \pm 0.4	54.2 \pm 2.2
BiC [59]	Coreset	53.7 \pm 0.4	87.5 \pm 0.9	45.9 \pm 1.8	81.9 \pm 2.0	37.5 \pm 3.2	71.7 \pm 3.4
Ours	Synthetic	43.9 \pm 0.9	78.6 \pm 1.1	33.7 \pm 1.2	69.6 \pm 1.6	20.0 \pm 1.4	52.5 \pm 2.5

order before sampling tasks, and do this with a consistent seed (different for each trial) so that results can be directly compared. We include supplementary details and metrics in our Appendix: additional results, additional experiment details, and hyperparameter selection.

Evaluation Metrics: Following prior works, we evaluate methods in the class-incremental learning setting using: (I) final performance, or the performance with respect to all past classes after having seen all N tasks (referred to as $A_{N,1:N}$); and (II) Ω , or the average (over all tasks) normalized task accuracy with respect to an offline oracle method [18]. We use index i to index tasks through time and index n to index tasks with respect to test/validation data (for example, $A_{i,n}$ describes the accuracy of our model after task i on task n data). Specifically:

$$A_{i,n} = \frac{1}{|\mathcal{D}_n^{test}|} \sum_{(x,y) \in \mathcal{D}_n^{test}} \mathbb{1}(\hat{y}(x, \theta_{i,n}) = y \mid \hat{y} \in \mathcal{T}_n) \quad (16)$$

$$\Omega = \frac{1}{N} \sum_{i=1}^N \sum_{n=1}^i \frac{|\mathcal{T}_n|}{|\mathcal{T}_{1:i}|} \frac{A_{i,1:n}}{A_{offline,1:n}} \quad (17)$$

where $A_{offline}$ is the task accuracy trained in the offline setting (i.e., the upper-bound performance). Ω is designed to evaluate the global task and is therefore computed with respect to all previous classes. For the final task accuracy in our results, we will denote $A_{N,1:N}$ as simply A_N .

Data-Free Class-Incremental Learning - CIFAR-100 Benchmark:

Our first benchmark is ten-task class-incremental learning on the CIFAR-100 dataset [29] which contains 100 classes of 32x32x3 images. Following prior work [59], we train with a 32-layer ResNet [19] for 250 epochs. The learning rate is set to 0.1 and is reduced

by 10 after 100, 150, and 200 epochs. We use a weight decay of 0.0002 and batch size of 128. Using a simple grid search to find hyperparameters for Eq (7), we found $\{\alpha_{con}, \alpha_{div}, \alpha_{stat}, \alpha_{prior}, \alpha_{temp}\}$ as $\{1, 1, 5e1, 1e-3, 1e3\}$ (these hyperparameters are not introduced by our method). For Eq. (15) in our method, we found $\{\lambda_{kd}, \lambda_{ft}\}$ as $\{1e-1, 1\}$, and we use prior reported loss-weighting hyperparameters for our implementations of other methods. We use a temperature scaling of 2 for all softmax knowledge distillation instances. We model \mathcal{F} with the same parameters as [39] and train Eq. (7) after starting each task using 5,000 training steps of Adam optimization (learning rate 0.001).

The results are given in Table 1. We see that our method outperforms not only the DFCIL methods (including a 25.1% increase in final task accuracy over DeepInversion), but even the generative approach (despite their use of significant additional memory between tasks). To our surprise, we found DGR [49] to perform poorly for class-incremental learning on this dataset (and in fact every dataset we experiment with); this finding is repeated in another recent work [54] which also found DGR to perform worse¹ than Base. We are not surprised to see that LwF [34] performs worse than naive rehearsal, as this is also common for class-incremental learning [55, 54]. Finally, we observe that synthetic data does not improve LwF’s performance. This is consistent with our finding that the feature embedding prioritizes domains over semantics when using standard distillation strategies.

¹We did not implement additional generative-replay results because that is not the focus of our paper. Instead, we compare to 1) other DFCIL methods to show our method performs best in our setting, and 2) SOTA replay-based methods to show our method performs close to SOTA despite not storing replay data.

Table 3: Ablation Study Results (%) for ten-task class-incremental learning on CIFAR-100. Results are reported as an average of 3 runs.

Metric (\uparrow)	A_N	Ω
Full Method	33.7 \pm 1.2	69.6 \pm 1.6
Ablate Task Balancing Loss Weighting [31, 49, 59]	23.4 \pm 1.5	62.2 \pm 2.4
Replace Modified CE Loss, Eq. (11), w/ Standard CE Loss	16.5 \pm 0.5	46.6 \pm 0.9
Ablate Real Data Distillation: Eq. (13) w/ X	15.9 \pm 2.1	58.9 \pm 3.2
Ablate Synthetic Data Distillation: Eq. (13) w/ \hat{X}	12.7 \pm 7.4	55.3 \pm 8.1
Ablate FT-CE Loss: Eq. (14)	9.8 \pm 0.6	35.9 \pm 1.3

Table 4: Results (%) for class-incremental learning on five task ImageNet-50. A coreset of 2000 images is leveraged for replay-based methods, and thus *these methods do not meet problem the DFCIL constraints*. Results are reported as a single run.

Method	Replay Data	A_N (\uparrow)
Upper Bound	None	89.8
LwF [34]	None	19.4
Naive Rehearsal	Coreset	78.9
LwF [34]	Coreset	84.8
Ours	Synthetic	71.5

Class-Incremental Learning with Replay Data - CIFAR-100 Benchmark:

In Table 2, we compare our method (which does *not store replay data*) to other methods which *do store replay data*. We use the same number of coreset images as the [6, 59]. We found that our method can perform *significantly greater than LwF and Rehearsal, which store replay data*. We also compare our method with a SOTA replay-based class-incremental learning method: Bias Correction (BiC) [59]. Despite not storing any replay data, our method performs roughly in between BiC and LwF, though there is still a considerable (and expected) gap between our data-free approach and BiC. In summary, these results indicate that *our method achieves State-of-the-Art performance for Data-Free Class-Incremental Learning*, and *our method closes much of the performance gap between Data-Free Class-Incremental Learning and State-of-the-Art replay-based methods*.

Ablation Study - CIFAR-100 Benchmark: We separate the components of our method to independently evaluate their effect on final performance, shown in Tables 3. We first look at the effect of removing task balancing loss weighting. As previously reported [31, 49, 59], this loss weighting has a significant effect on our performance. Next, we show that replacing the modified cross-entropy loss with standard cross-entropy loss cuts performance in half. Similarly, we show that ablating real and synthetic data distillation have the same effect. This indicates that 1) all three of these losses are crucial for our approach, and 2) we still establish SOTA performance despite removing any of these three losses. Finally, we see that removing the fine-tuning cross-entropy loss has the largest effect on performance. Conceptually, this makes sense because without this loss there is no way to distinguish new task classes from previous task classes.

Class-Incremental Learning with Replay Data - Im-

ageNet Benchmark: Finally, we use the ImageNet dataset [47] to demonstrate how our method performs on large scale 224x224x3 images. Following prior work [59], we train with a 18-layer ResNet [19] for 100 epochs. The learning rate is set to 0.1 and is reduced by 10 after 30, 60, 80, and 90 epochs. We use a weight decay of 0.0001 and batch size of 128. We use the same class-shuffling seed as prior work [44, 59] and report top-5 accuracy on ten tasks of 100 classes. We scale down to 20k coreset images used in the full ImageNet experiments to 2k, consistent with the relative number of classes. We also double the number of training steps used to train \mathcal{F} . Every other experiment detail is the same as the CIFAR-100 experiments.

The results are given in Table 4. Importantly, this experiment is significant because the number of parameters stored for replay (2000*224*224*3 = 3e8) *far exceeds* the number of parameters *temporarily* stored for synthesizing images (3.3e6). Despite requiring only 100 times fewer parameters to store, our method performs reasonably close to replay on this large-scale image experiment. We also far outperform LwF, which is the only DFCIL method to have been previously tried on large-scale ImageNet experiments. Additional experiments on the challenging Tiny-ImageNet dataset [30], which demonstrate the scalability of our method, are available in the Appendix.

8. Conclusions

We show that existing class-incremental learning methods perform poorly when learning a new task with real training data and preserving past knowledge with synthetic distillation data. We then contribute a new method which achieves SOTA performance for data-free class-incremental learning, and is comparable to SOTA replay-based approaches. Our research vision is to *eliminate the need for storing replay data* in class-incremental learning, enabling broad and practical applications of computer vision. An incremental learning solution which does not store data will provide immediate impact to computer vision applications such as reducing memory requirements for autonomous vehicles (which generate an inordinate amount of data), eliminating the need to transfer private medical data for medical imaging research collaboration (which is limited by strict legal protections), or removing the need to track private data for personal device user recommendation systems.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *ECCV*, 2018. [2](#)
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In *Advances in Neural Information Processing Systems*, pages 11849–11860, 2019. [2](#)
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pages 11816–11825, 2019. [2](#)
- [4] Yogesh Balaji, Mehrdad Farajtabar, Dong Yin, Alex Mott, and Ang Li. The effectiveness of memory replay in large scale continual learning. *arXiv preprint arXiv:2010.02418*, 2020. [3](#)
- [5] Brett K Beaulieu-Jones, Zhiwei Steven Wu, Chris Williams, Ran Lee, Sanjeev P Bhavnani, James Brian Byrd, and Casey S Greene. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122, 2019. [1](#)
- [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. [2](#), [8](#)
- [7] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-GEM. In *International Conference on Learning Representations*, 2019. [2](#)
- [8] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. Continual learning with tiny episodic memories. *arXiv preprint arXiv:1902.10486*, 2019. [2](#)
- [9] Hanting Chen, Yunhe Wang, Chang Xu, Zhaohui Yang, Chuanjian Liu, Boxin Shi, Chunjing Xu, Chao Xu, and Qi Tian. Data-free learning of student networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3514–3522, 2019. [3](#), [4](#)
- [10] Anupam Das, Martin Degeling, Xiaoyou Wang, Junjue Wang, Norman Sadeh, and Mahadev Satyanarayanan. Assisting users in a world full of cameras: A privacy-aware infrastructure for computer vision applications. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1387–1396. IEEE, 2017. [1](#)
- [11] Sayna Ebrahimi, Mohamed Elhoseiny, Trevor Darrell, and Marcus Rohrbach. Uncertainty-guided continual learning with bayesian neural networks. *arXiv preprint arXiv:1906.02425*, 2019. [2](#)
- [12] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. *arXiv preprint arXiv:2003.09553*, 2020. [2](#)
- [13] Alexander Gepperth and Cem Karaoguz. Incremental learning with self-organizing maps. *2017 12th International Workshop on Self-Organizing Maps and Learning Vector Quantization, Clustering and Data Visualization (WSOM)*, pages 1–8, 2017. [2](#)
- [14] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. [5](#)
- [15] Tian Guo. Cloud-based or on-device: An empirical study of mobile deep inference. In *2018 IEEE International Conference on Cloud Engineering (IC2E)*, pages 184–190. IEEE, 2018. [1](#)
- [16] Matan Haroush, Itay Hubara, Elad Hoffer, and Daniel Soudry. The knowledge within: Methods for data-free model compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8494–8502, 2020. [3](#), [4](#)
- [17] Tyler L Hayes, Nathan D Cahill, and Christopher Kanan. Memory efficient experience replay for streaming learning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9769–9776. IEEE, 2019. [2](#)
- [18] Tyler L Hayes and Christopher Kanan. Lifelong machine learning with deep streaming linear discriminant analysis. *arXiv preprint arXiv:1909.01520*, 2019. [7](#)
- [19] 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. [7](#), [8](#)
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [2](#)
- [21] 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](#)
- [22] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via re-balancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. [2](#), [3](#), [5](#)
- [23] Yen-Chang Hsu, Yen-Cheng Liu, Anita Ramasamy, and Zsolt Kira. Re-evaluating continual learning scenarios: A categorization and case for strong baselines. *arXiv preprint arXiv:1810.12488*, 2018. [3](#)
- [24] Xu Ji, Joao Henriques, Tinne Tuytelaars, and Andrea Vedaldi. Automatic recall machines: Internal replay, continual learning and the brain. *arXiv preprint arXiv:2006.12323*, 2020. [3](#)
- [25] Nitin Kamra, Umang Gupta, and Yan Liu. Deep generative dual memory network for continual learning. *arXiv preprint arXiv:1710.10368*, 2017. [1](#), [2](#)
- [26] Ronald Kemker and Christopher Kanan. Fearnnet: Brain-inspired model for incremental learning. *International Conference on Learning Representations (ICLR)*, 2018. [1](#), [2](#)
- [27] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. *AAAI Conference on Artificial Intelligence*, 2018. [2](#)
- [28] 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. [2](#)
- [29] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Tech Report*, 2009. [7](#)
- [30] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7, 2015. [8](#)
- [31] Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 312–321, 2019. [2](#), [3](#), [6](#), [8](#)
- [32] Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A neural dirichlet process mixture model for task-free continual learning. *arXiv preprint arXiv:2001.00689*, 2020. [2](#)
- [33] Timothée Lesort, Hugo Caselles-Dupré, Michael Garcia-Ortiz, Andrei Stoian, and David Filliat. Generative models from the perspective of continual learning. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019. [2](#)
- [34] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [35] Bo Liu, Xuesu Xiao, and Peter Stone. Lifelong navigation. *arXiv preprint arXiv:2007.14486*, 2020. [1](#)
- [36] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. *arXiv preprint arXiv:1705.03550*, 2017. [2](#)
- [37] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks. *arXiv preprint arXiv:1710.07535*, 2017. [3](#)
- [38] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pages 6470–6479, USA, 2017. Curran Associates Inc. [2](#)
- [39] Liangchen Luo, Mark Sandler, Zi Lin, Andrey Zhmoginov, and Andrew Howard. Large-scale generative data-free distillation. *arXiv preprint arXiv:2012.05578*, 2020. [3](#), [4](#), [7](#)
- [40] Davide Maltoni and Vincenzo Lomonaco. Continuous learning in single-incremental-task scenarios. *Neural Networks*, 116:56–73, 2019. [2](#)
- [41] Alexander Mordvintsev, Christopher Olah, and Mike Tyka. Inceptionism: Going deeper into neural networks, 2015. [3](#)
- [42] Vaishnavh Nagarajan, Colin Raffel, and Ian J Goodfellow. Theoretical insights into memorization in gans. In *Neural Information Processing Systems Workshop*, 2018. [1](#)
- [43] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jah-nichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11321–11329, 2019. [2](#)
- [44] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. icarl: Incremental classifier and representation learning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR’17*, pages 5533–5542, 2017. [2](#), [8](#)
- [45] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. [2](#)
- [46] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lilicrap, and Gregory Wayne. Experience replay for continual learning. In *Advances in Neural Information Processing Systems*, pages 348–358, 2019. [2](#)
- [47] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. [8](#)
- [48] 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. [2](#)
- [49] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2990–2999. Curran Associates, Inc., 2017. [1](#), [2](#), [6](#), [7](#), [8](#)
- [50] James Smith, Jonathan Balloch, Yen-Chang Hsu, and Zsolt Kira. Memory-efficient semi-supervised continual learning: The world is its own replay buffer. *arXiv preprint arXiv:2101.09536*, 2021. [2](#)
- [51] James Smith, Cameron Taylor, Seth Baer, and Constantine Dovrolis. Unsupervised progressive learning and the STAM architecture. *arXiv preprint arXiv:1904.02021*, 2021. [2](#)
- [52] Prathap Soma, Ravi Kumar Jatoth, and Hathiram Nenavath. Fast and memory efficient de-hazing technique for real-time computer vision applications. *SN Applied Sciences*, 2(3):1–10, 2020. [1](#)
- [53] Michalis K Titsias, Jonathan Schwarz, Alexander G de G Matthews, Razvan Pascanu, and Yee Whye Teh. Functional regularisation for continual learning with gaussian processes. In *International Conference on Learning Representations*, 2019. [2](#)
- [54] Gido M van de Ven, Hava T Siegelmann, and Andreas S Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, 11(1):1–14, 2020. [2](#), [7](#)
- [55] Gido M van de Ven and Andreas S Tolias. Generative replay with feedback connections as a general strategy for continual learning. *arXiv preprint arXiv:1809.10635*, 2018. [2](#), [7](#)
- [56] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. [6](#)
- [57] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. *arXiv preprint arXiv:1906.00695*, 2019. [2](#)
- [58] Liyuan Wang, Bo Lei, Qian Li, Hang Su, Jun Zhu, and Yi Zhong. Triple memory networks: a brain-inspired method for continual learning. *arXiv preprint arXiv:2003.03143*, 2020. [1](#)
- [59] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. [3](#), [6](#), [7](#), [8](#)
- [60] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-

inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8724, 2020. [1](#), [3](#), [4](#), [5](#), [6](#), [7](#)

- [61] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, 2017. [2](#)
- [62] Yuqing Zhu, Xiang Yu, Manmohan Chandraker, and Yu-Xiang Wang. Private-knn: Practical differential privacy for computer vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#)