

Efficient Expansion and Gradient Based Task Inference for Replay Free Incremental Learning

Soumya Roy
Amazon

meetsoumyaroy@gmail.com

Vinay Verma
Amazon/Duke University

vinayugc@gmail.com

Deepak Gupta
Amazon

deepakgupta.cbs@gmail.com

Abstract

This paper proposes a simple but highly efficient expansion-based model for continual learning. The recent feature transformation, masking and factorization-based methods are efficient, but they grow the model only over the global or shared parameter. Therefore, these approaches do not fully utilize the previously learned information because the same task-specific parameter forgets the earlier knowledge. Thus, these approaches show limited transfer learning ability. Moreover, most of these models have constant parameter growth for all tasks, irrespective of the task complexity. Our work proposes a simple filter and channel expansion-based method that grows the model over the previous task parameters and not just over the global parameter. Therefore, it fully utilizes all the previously learned information without forgetting, which results in better knowledge transfer. The growth rate in our proposed model is a function of task complexity; therefore for a simple task, the model has a smaller parameter growth while for complex tasks, the model requires more parameters to adapt to the current task. Recent expansion-based models show promising results for task incremental learning (TIL). However, for class incremental learning (CIL), prediction of task id is a crucial challenge; hence, their results degrade rapidly as the number of tasks increase. In this work, we propose a robust task prediction method that leverages entropy weighted data augmentations and the model's gradient using pseudo labels. We evaluate our model on various datasets and architectures in the TIL, CIL and generative continual learning settings. The proposed approach shows state-of-the-art results in all these settings. Our extensive ablation studies show the efficacy of the proposed components.

1. Introduction

Recent deep learning models outperform humans on many challenging tasks in a static environment [8, 37]. However, in a continuously changing environment where novel tasks arrive sequentially, humans significantly outper-

form deep learning models. When presented with sequential tasks, the model suffers from catastrophic forgetting and remembers only the current task sequence. Continual learning refers to continuously learning and adapting to new environments while exploiting knowledge acquired from the past tasks without forgetting.

The TIL and CIL are two most popular settings in continual learning. In the TIL setting, task ids are known during training and inference; however, in CIL, task ids are present only during training. The recent literature leverages three approaches to solve the continual learning problem. The replay-based methods [30, 31, 50] are the most popular approaches and show promising results, but we have to store a fraction of previous task samples for replay. Using past samples may violate privacy and increase training and storage costs. Moreover, learning is biased towards the current task since replay samples from the past tasks are limited. Regularization-based methods [1, 4, 13] regularize the previous task parameters to overcome catastrophic forgetting while learning the current task. They provide sub-optimal solutions and work well only for a limited task sequence. The expansion-based methods [33, 34, 51, 52, 58] are promising, as they can model many task sequences, but efficient expansion and task prediction are the critical bottlenecks. Therefore, most of the models work only for the simple TIL setting, where during inference, task ids are available. Recently [39, 44, 47] propose efficient expansion-based models, but their models do not fully utilize the previously learned parameters. They fine-tune their expansion parameter for every new task; therefore, it does not remember previous information and only the current task knowledge transfers to the next task. So, their knowledge transfer is not optimal. Moreover, most of these approaches assume that all tasks have same complexity; therefore, their expansion parameters grow equally. However, in practice, tasks are diverse and can have different complexities. Task prediction is another key challenge in expansion-based approaches; it is sensitive to the parameter and its performance decreases rapidly as the number of tasks increase.

In this work, we propose a simple and highly efficient

filter and channel expansion-based model which is generic and can be applied to any convolutional network. The model fully utilizes the previously learned information, hence providing better knowledge transfer. Our expansion-based model learns a task-specific parameter for each novel task. However, during learning, the i^{th} task model uses the global parameter and all learned task-specific parameters from 2 to the $(i - 1)^{th}$ task. Parameters up to the $(i - 1)^{th}$ task preserve all the previously learned information and maximize knowledge transfer during learning of the current task. The expansion in the task-specific parameter is achieved by increasing the number of filters in each layer. We also propose a gradient-based approach to grow the task-specific parameter as a function of the task complexity; therefore for a simple task, the model has a smaller parameter growth while for complex tasks, the model requires more parameters to adapt to the current task. Our expansion-based model shows promising results in the TIL scenario. We further propose a replay-free task prediction method to solve the more challenging CIL setting. The task prediction method leverages entropy weighted augmentations and pseudo label to calculate the sample loss. Finally, we approximate the model’s gradient by the mean gradient of each layer and the norm of the task gradient is used to predict the task id.

2. Related Work

Recently there has been a lot of interest in the continual learning paradigm [27], because of its extensive applicability to various domains. We can broadly divide modern continual learning approaches into three categories - Replay, Regularization and Expansion-based models. Replay-based methods [3, 29–32, 41, 48] keep a memory bank to store a fraction of samples from previous tasks. Handling sample bias [30] is a crucial challenge since the current task has a large number of samples, while the previous task samples are very few. Generative replay methods [21, 36, 42] learn a generative model for the sample replay. Regularization [1, 13, 19, 20] is another popular approach that regularizes the weight learned during previous tasks while training on the current task. SI [55], EWC [13] and IMM [19] focus on regularizing previous task weight based on their importance. These approaches cannot model a large number of tasks and model performance degrades quickly as the tasks increase. They provide sub-optimal solutions since they try to learn a joint weight that can generalize across all tasks. Recent methods like IL2A [61], SSRE [62] and FeTrIL [28] uses class prototypes to tackle replay-free class-incremental learning and are baselines for our method. TCL [12] provides a theoretical justification for decomposing the CIL problem into two sub-problems - TIL and task prediction. Task prediction is done using supervised contrastive learning, ensemble class prediction and output calibration by

leveraging replay data.

The expansion-based models are closely related to our work, hence we mostly focus on these approaches. Recently, a wide range of methods [6, 23, 24, 33–35, 39, 44, 45, 49, 51, 52, 58] propose expansion-based strategies to incorporate the growing task sequence. PNNs [34] directly adds a new neural network column for each novel task. These approaches freeze the weights of the previous tasks to overcome catastrophic forgetting and lateral connections help forward knowledge transfer. Side-tuning [58] proposes a simpler approach by training a small task-specific network and fusing the output to the base network. DEN [52] not only focuses on network expansion but also optimizes sub-problems like selective training, dynamic model expansion using loss threshold and duplication. RCL [49] uses reinforcement learning to determine the growth of the architecture. Recent advancements leverage Bayesian non-parametric models where the data itself determines the expansion [16, 18, 25], but these methods work well for small datasets and few task sequences like MNIST and CIFAR-10. EFT [44] partitions a model into global and task-specific local parameters and leverages efficient convolution operations to construct these local transforms. The global network can be any architecture and the method outperforms most baselines on diverse task sequences in both TIL and CIL settings.

Masking [22–24, 35, 47] is another expansion-based strategy that learns different binary/ternary masks per task. Iterative training and pruning strategies [11, 22] have also been proposed for expansion-based continual learning. These approaches are costly to train, since pruning is an expensive step. PackNet [23] requires saving masks to recover networks of previous models, which can take lots of storage space as the number of tasks grow. HAT [35] proposes hard attention masks for each task. TFM [24] applies ternary masks to feature maps, which result in less memory per mask, as the feature maps are often smaller than the number of weights in model. TFM uses all previous tasks to learn the mask of the current task; thus, it has similar motivations as our method. APD [51] decomposes the network parameters into task-shared and sparse task-specific parameters; however, the significant changes made to architecture make it harder to scale and prevent the use of pre-trained weights. SupSup [47] finds a super-mask for each task and uses gradient-based optimization to speed up inference; also, the super masks can be stored in a fixed-size Hopfield network [10]. Masking approaches are promising, but they are mostly limited to the simpler TIL setting. Choosing a suitable mask requires predicting the task id, which is challenging. Our proposed approach shows promising results in both scenarios without relying on replay samples.

3. Proposed Model

In this section, we provide a detailed description of our proposed model. We use three key components to overcome catastrophic forgetting.

3.1. Notations

In incremental learning, tasks $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_T$ arrive sequentially. Each task $\mathcal{T}_i = \{(\mathbf{x}_k, y_k)\}_{k=1}^{N_i}$ where \mathbf{x}_k has label $y_k \in \mathcal{Y}_i$ and \mathcal{Y}_i is the set of all classes in task \mathcal{T}_i ($|\mathcal{Y}_i| = K_i$). During training on the i^{th} task, only \mathcal{T}_i is available to train the model and $\mathcal{Y}_i \cap \mathcal{Y}_j = \phi$. Our proposed model can be applied to TIL as well as the CIL setting. During training and inference of the task \mathcal{T}_i , the task id i is known in the TIL setting. However, in the CIL setting, task id i is unknown during inference.

3.2. Efficient Dynamic Expansion (ablated in Sec. 5.3)

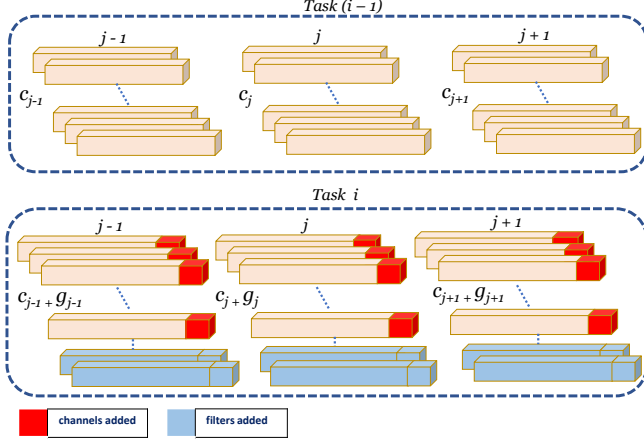


Figure 1. Expansion of the j^{th} layer during task i after training on the $(i-1)^{\text{th}}$ task.

In this work, we propose an efficient, dynamically expandable model which grows with the number of tasks and accumulates all the previously learned knowledge. Let \mathcal{M}_{θ_i} be the model for task \mathcal{T}_i where θ_i is the model parameter. The model parameter $\theta_i = \{\Omega, \tau_i, \tau_{2:i-1}\}$, i.e., it contains three types of parameters - global parameter (Ω), task-specific parameter (τ_i) and all the learned task-specific parameters before the current task ($\tau_{2:i-1}$). Therefore, the parameter θ is growing with each novel task sequence and $|\theta_1| < |\theta_2| < \dots < |\theta_T|$ where $|\cdot|$ represents cardinality.

Now we will describe our expansion method for a particular task \mathcal{T}_i . For task \mathcal{T}_{i-1} , we can represent the j^{th} convolutional layer for task model $\mathcal{M}_{\theta_{i-1}}$ as $l_{i-1}^j \in \mathbb{R}^{c_j \times k \times k \times d_j}$ where c_j is the number of filters, k is the kernel size and d_j is the number of feature maps of layer $(j-1)$. After training on the $(i-1)^{\text{th}}$ task, we freeze its parameter θ_{i-1} . Let us assume that g_{j-1} , g_j and g_{j+1} are the *growth*

rates for the layers $(j-1)$, j and $(j+1)$ respectively for the next task \mathcal{T}_i . Before training the model for task \mathcal{T}_i , we will grow the model by the defined growth rates and this growth is local to the i^{th} task (and hence the name task-specific parameter). The i^{th} task model adds g_{j-1} , g_j and g_{j+1} number of filters at the $(j-1)$, j and $(j+1)$ layers respectively. Therefore, for the i^{th} task at the j^{th} layer, we have $c_j + g_j$ number of filters, but layer $(j-1)$ will also produce $c_{j-1} + g_{j-1}$ feature maps. To accommodate these feature maps, we have to increase the number of channels in each filter of layer j . So $l_i^j \in \mathbb{R}^{(c_j + g_j) \times k \times k \times d_j'}$ where $d_j' = c_{j-1} + g_{j-1}$ is the number of feature maps after adding the task-specific filters to layer $(j-1)$. Similarly, the next layer $l_i^{j+1} \in \mathbb{R}^{(c_{j+1} + g_{j+1}) \times k \times k \times d_{j+1}'}$, where $d_{j+1}' = c_j + g_j$. Figure 1 shows our expansion strategy at the particular layer j during training of task \mathcal{T}_i . For each task, we train batch norm and final linear layer from scratch. So the global parameter (Ω) is the parameter of the first task model without batch norm and final linear layer.

3.3. Gradient aggregation for task prediction (ablated in Sec. 5.4)

The proposed expansion-based model shows promising results in the TIL setting. However, in the CIL setting, we don't know the task id during inference; hence we have to predict it. In this work, we propose a robust method for task prediction that can predict the task id for a large number of task sequences.

Let $\mathcal{M}_{\theta_1}, \mathcal{M}_{\theta_2}, \dots, \mathcal{M}_{\theta_T}$ be the learned models for the tasks $\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_T$ with parameters $\theta_1, \theta_2, \dots, \theta_T$. Once we have finished training for the task \mathcal{T}_i , we can predict the sample's task id of any classes till the i^{th} task. The gradient embedding of test sample \mathbf{x}_k , for the model with parameter θ_i over the loss function $\mathcal{L}(\mathbf{x}_k, \theta_i)$, can be computed as:

$$\theta'_{ik} = \nabla_{\theta_i} \mathcal{L}(\mathbf{x}_k, \theta_i) \quad (1)$$

To predict the task id of sample \mathbf{x}_k , we measure the gradient norm using each model as follows:

$$\hat{i} = \arg \min_{i \in T} \frac{\|\theta'_{ik}\|}{|\theta'_{ik}|} \quad (2)$$

Here $|\theta'_{ik}|$ is the number of parameters in the gradient vector. We need to normalize the embedding, as each task has a different gradient length and the norm of a longer embedding can be higher. Once we have task id \hat{i} for the sample \mathbf{x}_k , we can choose the model $\mathcal{M}_{\theta_{\hat{i}}}$ to evaluate the sample.

Minimizing the norm of loss gradient leads to flatter minima in weight space [59]. A flat minima is a large connected region in weight space where the loss remains approximately constant and it is optimal minima for generalization [9, 59]. In Eq. 2, we test the shape of the loss function \mathcal{L} using \mathbf{x}_k to estimate which model might perform better on that sample. We use ℓ_1 -norm as it gives marginally




Figure 2. Overview of the gradient aggregation method that leverages weighted cross-entropy using pseudo label.

better results than ℓ_2 -norm. For a fair comparison with earlier baselines, we do not use the gradient norm loss during training time.

For a convolutional network with L layers, we define the (flattened) gradient vector for layer j as $\theta'_{ik}{}^j$ ($j = 1 \dots L$). An advantage of considering layer-wise gradient is that it allows us to use task and layer dependant gradient modification functions. Such functions can help us give more *attention* to certain parameters in a layer. Thus we can rewrite the gradient embedding for task \mathcal{T}_i and input \mathbf{x}_k as the vector:

$$\theta'_{ik} = \text{CONCAT}(\mathcal{G}_i^1(\theta'_{ik}{}^1), \mathcal{G}_i^2(\theta'_{ik}{}^2), \dots, \mathcal{G}_i^L(\theta'_{ik}{}^L)) \quad (3)$$

where \mathcal{G}_i^j is the gradient modification function for task i and layer j and CONCAT concatenates multiple gradient vectors to create the final embedding vector. As our inference time grows linearly with the number of tasks, we can borrow the superposition idea of [47] to get sub-linear run times. However, in this paper, we do not focus on this direction. We present an overview of our entire method in Figure 2. Next we choose $\mathcal{L}(\mathbf{x}_k, \theta_i)$ and \mathcal{G}_i^j .

3.3.1 Choosing loss function

The most common loss function for training deep networks is cross-entropy. However, task prediction using a single sample \mathbf{x}_k may not be robust. So our approach leverages simple test-time data augmentations to improve task prediction performance. From input \mathbf{x}_k , we create a batch of A augmented samples $\mathcal{X}_k = [\mathbf{x}_k^1, \mathbf{x}_k^2 \dots, \mathbf{x}_k^A]$. For each augmented sample in \mathcal{X}_k , we consider the class with the maximum probability as it's pseudo label and add it to the array $\hat{\mathcal{Y}}_k$. Subsequently, we take the maximally occurring pseudo label in $\hat{\mathcal{Y}}_k$ as the batch pseudo label \hat{y}_k . Mathematically, this can be expressed as:

$$\begin{aligned} \hat{\mathcal{Y}}_k &= [\arg \max_{K_i} \theta_i(\mathbf{x}_k^1), \dots, \arg \max_{K_i} \theta_i(\mathbf{x}_k^A)] \\ \hat{y}_k &= \text{MAXCOUNT}(\hat{\mathcal{Y}}_k) \end{aligned} \quad (4)$$

where MAXCOUNT takes an array as input and returns the maximally occurring element in the array. $\theta_i(\mathbf{x}_k^1)$ is the output of network θ_i for image \mathbf{x}_k^1 .

The pseudo label \hat{y}_k is used to calculate the cross-entropy loss over the batch \mathcal{X}_k using the model θ_i . This loss measures the robustness of the model with respect to sample perturbations. However, this augmentation strategy can create samples of different complexities, i.e., some samples can be more *confusing* to the model than others. So, instead of giving equal weights to every augmented sample, we calculate the batch cross-entropy loss in a weighted manner. We leverage the entropy measure to estimate the uncertainty of each augmented sample. For augmented sample \mathbf{x}_k^a and model θ_i , we use $\mathcal{CE}(\theta_i(\mathbf{x}_k^a), \hat{y}_k)$ and $\mathcal{ENT}(\theta_i(\mathbf{x}_k^a))$ to denote the sample cross-entropy loss and sample entropy respectively. So the final loss function can be expressed as:

$$\mathcal{L}(\mathbf{x}_k, \theta_i) = \frac{1}{A} \sum_{a \in A} \mathcal{CE}(\theta_i(\mathbf{x}_k^a), \hat{y}_k) \times \mathcal{ENT}(\theta_i(\mathbf{x}_k^a)) \quad (5)$$

We use this loss to calculate the gradient in Eq. 1. It should be noted that calculating the gradient of Eq. 5 for an augmentation \mathbf{x}_k^a is mathematically equivalent to weighing the gradients of KL-divergence and entropy terms in the corresponding cross-entropy loss $\mathcal{CE}(\theta_i(\mathbf{x}_k^a), \hat{y}_k)$.

3.3.2 Choosing gradient modification function

The cardinality of the gradient in Eq. 1 is equal to the number of model parameters. Therefore, any operation on this vector is very costly. Moreover, noise in gradient directions can decrease the model's robustness. Hence, we only use the mean gradient for each convolutional filter. For fully

connected layers, we use the mean gradient for each output neuron. In the standard convolutional network, the initial layers capture generic information about the input image while the later layers contain task-specific information. Hence, we use the gradients of the last two convolutional layers and the fully connected layer. We call this strategy *mean filters*; it reduces the gradient cardinality by around 99.98% and even helps improve task prediction.

3.4. Static and Adaptive Growth Rate (ablated in Sec. 5.1)

Most expansion-based models [39, 44] consider the growth rate as a hyperparameter and irrespective of the task complexity, it remains constant across all tasks. Here, we propose an adaptive growth rate for the expansion parameter that depends on the task complexity. The calculation of task complexity, without availability of the previous task samples, is the key challenge. For the i^{th} task, we leverage the current task samples \mathcal{T}_i , the $(i-1)^{th}$ task model and the mean gradient on the previous task samples \mathcal{T}_{i-1} to construct a compatibility score. Subsequently, we use this score to decide the growth rate for the i^{th} task.

Let’s assume that before training for task \mathcal{T}_i , we expand each layer j of the $(i-1)^{th}$ task model by g_j filters (i.e., growth rate is g_j). Using scaling factor $\alpha \in [0, 1]$, we can write g_j as:

$$g_j = [\alpha \times g_j^{min} + (1 - \alpha) \times g_j^{max}] \quad (6)$$

where g_j^{max} and g_j^{min} are respectively the pre-defined maximum and minimum growth rates for layer j and $[\cdot]$ is the rounding function.

In case of static growth, the growth rate for a layer is constant for all tasks and $\alpha = 0$. However, in adaptive growth, we expand our model based on the relative complexity of the novel task \mathcal{T}_i and use α to encode how *similar* \mathcal{T}_i is with \mathcal{T}_{i-1} . Since our models are optimized using gradient-based methods, we can measure the uncertainty of a sample using gradients [2]. Given the $(i-1)^{th}$ task model, we can compare the model confidences in two samples \mathbf{x}_1 and \mathbf{x}_2 using $\theta'_{(i-1)1}$ and $\theta'_{(i-1)2}$ from Eq. 1. For tasks \mathcal{T}_{i-1} and \mathcal{T}_i , we can write the scaling factor α as:

$$\alpha = |\mathcal{MEAN}(\{\theta'_{(i-1)k}\}_{k \in \mathcal{T}_i}) \cdot \mathcal{MEAN}(\{\theta'_{(i-1)k}\}_{k \in \mathcal{T}_{i-1}})| \quad (7)$$

where \mathcal{MEAN} gives us the normalized mean vector from a set of input vectors. We can use this α to decide the growth rate for the i^{th} task using Eq. 6. So for each task, we only need to remember the mean gradient of the last task. We can reduce the gradient cardinality to around 0.02% of its original size using mean filters strategy. Based on the type of growth, we get two model variants - static parameter growth (SPG) model and adaptive parameter growth (APG) model. For a fair comparison between SPG and APG models, we set g_j^{min} as 1 and g_j^{max} as the growth rate used in SPG. The parameter growth rate is ablated in Sec. 5.2.

4. Experiments and Results

We conduct our experiments in TIL, CIL as well as generative continual learning settings. We also evaluate our method using different base architectures on diverse datasets. Our approach offers significant gains over the latest baselines in all the experiments.

4.1. Datasets and Architectures

Our experiments are performed on three datasets - CIFAR-100 [14], ImageNet-100 [15] and TinyImageNet [17]. We divide the CIFAR-100 dataset into 5, 10, and 20 task sequences (let’s call them CIFAR100/5, CIFAR100/10 and CIFAR100/20 respectively). CIFAR100/5 has a smaller task sequence, but each task contains 20 classes. On the other hand, CIFAR100/20 contains only 5 classes per task but has a large number of tasks. We evaluate our method on this dataset using ResNet-18 [8] architecture for CIFAR datasets. For the ImageNet-100 dataset, we use the same class subset and class order as DER [50]. Similar to CIFAR-100, we divide the ImageNet-100 dataset into 5, 10, and 20 tasks. The ImageNet-100 dataset is challenging as it contains relatively bigger images, i.e., the image size is 224×224 . We use the standard ResNet-18 architecture for experimenting on this dataset. Tiny ImageNet is a smaller subset of the ImageNet [15] dataset, contains 200 classes with a 64×64 resolution. To demonstrate the efficacy of our proposed model across different architectures, we use the VGG-16 [38] architecture with batch norm for the Tiny ImageNet dataset. We divide the dataset into 10 tasks and evaluate the model performance in the TIL setting. Furthermore, we explore our continual learning approach using the StackGAN-v2 [56] architecture for three diverse tasks - cats (ImageNet [5]), birds (CUB-200 [46]) and churches (LSUN [53]).

4.2. Baselines

Our proposed method leverages network expansion to learn new tasks and does not use any pretrained model or sample storage for replay. Therefore, comparison with replay-based methods is not fair. So we compare our model with both regularization and expansion-based approaches. We use various regularization based methods like LwF [20], EWC [13], SI [55], MAS [1], SDC [54], DMC [57], L2T [51], IL2A [61], LwF+adBiC [40], SSRE [62] and FeTrIL [28] as our baselines. We also use expansion-based approaches like PNN [34], HAT [35], PB [22], PackNet [23], TFM [24], APD [51] and EFT [44] for comparison. For the continual GAN experiment, we consider EWC [13], MeRGAN-RA [48] and EFT [44] as baselines. As we share our hyperparameters with EFT [44], we borrow most baseline results from that paper; the rest are obtained by either using the paper’s publicly available code or by running the PyCIL [60] framework for that model.




Figure 3. The comparison for the TIL scenario on TinyImageNet-200/10. The result for the i^{th} task is reported after learning all the tasks.

4.3. Implementation Details

We provide the model architecture, hyperparameter settings and other experimental details for CIFAR-100, ImageNet-100 and Tiny ImageNet datasets in the supplementary material. Moreover, in supplementary, we provide results for multiple seed values along with their standard deviations.

In the following section, we discuss the results in the TIL, CIL and generative continual learning settings. We use average incremental accuracy as the evaluation metric. For TIL, we compute it as the average accuracy over all tasks, including the initial one [44]. In case of CIL, we report the average accuracy for all seen classes [44, 61].

4.4. Results

4.4.1 Task Incremental Learning (TIL)

TIL is a relatively simpler setting where the task id is known during inference. Once we know the task id, we can select the model corresponding to that id for inference. We evaluate the TIL scenario for the Tiny ImageNet [17] dataset using the VGG-16 architecture. The Tiny ImageNet dataset contains 200 classes and we divide the dataset into 10 tasks each containing 20 classes. Our static parameter growth (SPG) model has an average parameter growth of 3.7% and achieves an average incremental accuracy of 68.6%. In comparison, EFT [44] has an average incremental accuracy of 66.8% with 3.6% average parameter growth. On the other hand, our adaptive parameter growth (APG) model gives an average incremental accuracy of 68.5% with a smaller average parameter growth of 3.4%. The task-wise accuracy for SPG model is shown in Fig. 3.

4.4.2 Class Incremental Learning (CIL)

The TIL setting requires us to know the task id during inference. This setting is not practical as, in the real world, we may not know the actual source of a datapoint. Thus, in expansion-based methods, prediction of the task id is a

Methods	CIFAR-100			ImageNet-100		
	5	10	20	5	10	20
LwF [20]	34.7	23.9	14.2	55.1	37.7	22.0
EWC [13]	25.0	16.4	9.6	-	-	-
SI [55]	29.6	23.3	13.3	-	-	-
MAS [1]	24.4	15.4	10.6	-	-	-
RWalk [4]	31.6	17.9	11.0	-	-	-
EWC+SDC [54]	-	19.3	-	-	-	-
DMC [57]	46.6	36.2	23.9	-	-	-
IL2A [61]	50.2	29.7	14.7	40.6	21.7	9.7
EFT [44]	52.7	45.5	30.3	49.2	42.5	36.3
LwF+adBiC [40]	44.1	33.3	19.5	-	-	-
SSRE [62]	41.9	30.0	13.7	37.5	24.0	16.2
FeTrIL [28]	45.0	34.3	20.8	44.0	31.2	19.8
Ours (SPG) [4.2%]	59.4	50.6	35.6	61.7	48.6	38.3
Ours (APG) [3.7%]	59.2	50.5	36.4	61.8	49.8	38.5

Table 1. Average incremental accuracy till last task for CIFAR-100 and ImageNet-100 in CIL setting. [X%] shows the average parameter growth of the model over all task sequences and datasets.

key challenge. To demonstrate the efficacy of our proposed gradient aggregation method for task prediction, we evaluate our model for diverse task sequences using average incremental accuracy. In all scenarios, our proposed model significantly improves over the previous state-of-the-art approaches.

In Table 1, we have shown results for 5, 10 and 20 task sequences on the CIFAR-100 dataset. Using 4.3% average parameter growth and static parameter growth (SPG) model variant, we achieve an absolute gain of 5.1% over our best baseline EFT [44] in the CIFAR100/10 setting. Similarly, using 4% average parameter growth, we get an absolute gain of 6.7% over EFT in the CIFAR100/5 setting. The CIFAR100/20 setting is the most challenging CIL scenario as the task prediction accuracy rapidly drops with increase in task length. However, even in CIFAR100/20 setting, we achieve an absolute gain of 5.3% over EFT using 4.1% average parameter growth. The adaptive parameter growth (APG) model variant shows similar average incremental accuracy, but has 12.2% less average parameter growth. Table 1 also shows results for 5, 10 and 20 task sequences on the ImageNet-100 dataset. Occasionally APG outperforms SPG as training large networks on small tasks can lead to overfitting [24].

Like [44, 61], we report the average accuracy for all seen classes. Hence, our results for recent exemplar-free class-incremental learning (EFCIL) methods like SSRE [62] and FeTrIL [28] differ from the results reported in such papers as they measure average accuracy for all states. In the supplementary material, we report results using the EFCIL experimental settings and evaluation metric.

Task i	Cats			Birds			Churches			Final Average
	1	2	3	1	2	3	1	2	3	
Finetune	29.0	156.9	189.6	-	21.2	174.5	-	-	11.4	125.2
EWC [13]	29.0	147.3	190.7	-	65.9	165.4	-	-	38.2	131.4
MeRGAN-RA [48]	29.0	56.4	58.2	-	50.9	53.7	-	-	23.2	45.1
EFT [44]	29.0	29.0	29.0	-	44.1	44.1	-	-	32.3	35.1
Ours	29.0	29.0	29.0	-	40.9	40.9	-	-	29.3	33.1

Table 2. Results for GAN (StackGAN-v2) when trained in sequential manner. FID is reported after training the final task.

4.4.3 Generative Continual Learning

Generative Adversarial Network (GAN) [7] also suffers from catastrophic forgetting [43] if data arrives as a sequence of tasks. However, training a separate model for each task is costly. Hence, we apply our proposed static filter expansion approach to train the GAN architecture in a continual learning fashion. We choose the StackGAN-v2 [56] architecture for the generative learning experiment the Table 2 shows the results for the continual generative models.

4.4.4 Heterogeneous Task Sequence

We evaluate our proposed static filter expansion approach on a heterogeneous task sequence. We have selected SVHN [26], CIFAR-10 and CIFAR-100 datasets for constructing a task sequence and perform our experiments in the increasing (SVHN→CIFAR10→CIFAR100) as well as decreasing (CIFAR100→CIFAR10→SVHN) orders of task complexity. The task sequence is heterogeneous in two ways - firstly, SVHN and CIFAR datasets are completely different in nature and secondly, the number of classes change (increase/decrease) between tasks, i.e., go from 10 → 100 or 100 → 10. The results for the heterogeneous setting are shown in Table 3. We can observe that in both the scenarios (i.e., forward and reverse sequence), our approach outperforms the recent baselines using just 4.7% average parameter growth.

5. Ablations

5.1. Adaptive Growth Rate: Toy Experiment

The CIFAR-100 [14] dataset contains 20 superclasses where each superclass has 5 similar classes. For this toy experiment, we consider 4 superclasses - *aquatic mammals*, *fish*, *vehicles 1* and *vehicles 2*. Using these superclasses, we construct two task sequences - *ordered* and *mixed*. Each task sequence has two tasks and each task has 10 distinct classes. Task 1 of the *ordered* sequence contains ($\{\textit{aquatic mammals}\} \cup \{\textit{fish}\}$) classes, while task 2 has ($\{\textit{vehicles 1}\} \cup \{\textit{vehicles 2}\}$) classes. On the other hand, the *mixed* sequence picks 2 – 3 classes from each superclass and constructs two tasks, each with 10 distinct classes. So the *ordered* task sequence has two very different tasks,

while the *mixed* task sequence has two somewhat similar tasks. In this experiment, after training on the first task, we find the gradient similarity between the first and second tasks to decide the growth rate for the second task. For *ordered* and *mixed* task sequences, we compute the gradient similarities $\alpha_{ordered}$ and α_{mixed} using Eq. 7. We observe $\alpha_{mixed} > \alpha_{ordered}$ and $(\alpha_{mixed} - \alpha_{ordered}) = 0.33$. This result confirms that our gradient measure can be used for finding similarities between tasks.

5.2. Parameter growth vs Performance

We define the parameter growth from task \mathcal{T}_i to \mathcal{T}_{i+1} as follows: $\frac{P_{i+1} - P_i + E_{i+1}}{P_i}$ where P_i is the total number of parameters used during task \mathcal{T}_i and E_{i+1} is the number of parameters in task \mathcal{T}_{i+1} which are trained from scratch during task \mathcal{T}_{i+1} (let’s call them *exclusive* parameters). For example, batch norm and the final linear layer in the ResNet-18 architecture are trained from scratch for every new task in our proposed approach and are thus counted as *exclusive* parameters. We define average parameter growth for T tasks as the mean of all T parameter growths. Since methods like EFT [44] grow their networks from the first task, we define the parameter growth for the first task as the change of parameters from an equivalent standard network.

In Table 4, we compare both our variants with our best baseline EFT [44] on 5, 10 and 20 splits of CIFAR-100. We show that our variants achieve significantly better results without any significant increase in average parameter growth. In Table 5, we show that the variants have almost no degradation in accuracy even when we significantly reduce average parameter growth. This is due to the forward transfer from previous tasks. We also show the total number of parameters that were used during training of all 10 tasks.

5.3. Forward Transfer

We have also performed experiments to demonstrate the forward transfer ability of our proposed filter expansion model. In this experiment, we have explored the scenario where our model grows over just the global or shared parameter. This type of growth is used in masking [47, 51] and EFT [44]. We have also shown the result for our current approach where the parameter grows over the previous tasks to accumulate all the earlier knowledge. The result for the CIFAR100/10 setting is shown in Fig. 4. Our approach

Task order	L2T		PB		PNN		APD		EFT		Ours	
	↓	↑	↓	↑	↓	↑	↓	↑	↓	↑	↓	↑
SVHN	10.7	88.4	96.8	96.4	96.8	96.2	96.8	96.8	96.8	95.5	96.6	95.6
CIFAR-10	41.4	35.8	83.6	90.8	85.8	87.7	90.1	91.0	89.2	90.4	90.5	91.7
CIFAR-100	29.6	12.2	41.2	67.2	41.6	67.2	61.1	67.2	64.6	71.5	64.5	71.8
Average	27.2	45.5	73.9	84.8	74.7	83.7	83.0	85.0	83.5	85.8	83.9	86.4

Table 3. Accuracy on the heterogeneous dataset sequence. ↓ follow the SVHN → CIFAR-10 → CIFAR-100 task order and ↑ represents CIFAR-100 → CIFAR-10 → SVHN task order.

Methods	5	10	20
EFT	52.7 (3.9%)	45.5 (3.9%)	30.3 (3.9%)
Ours (SPG)	59.4 (4.0%)	50.6 (4.3%)	35.6 (4.1%)
Ours (APG)	59.2 (3.2%)	50.5 (3.8%)	36.4 (3.7%)

Table 4. Average incremental accuracy till last task with average parameter growth in brackets for CIFAR-100.

Methods	Growth	Total Params	Accuracy
Baseline	-	11.2M	-
EFT	3.9%	15.8M	45.5
SPG	3.6%	15.8M	51.0
SPG	3.9%	16.3M	50.3
SPG	4.3%	17.0M	50.6
APG	3.1%	15.2M	49.8
APG	3.5%	15.7M	50.1
APG	3.8%	16.2M	50.5

Table 5. Average incremental accuracy with different average parameter growths for CIFAR-100/10. Total Params is the total number of parameters that were used for training all 10 tasks.




Figure 4. The effect of forward transfer; blue is for model which leverages previous task and yellow represents model which uses global parameter.

shows consistently better results across all tasks.

5.4. Task Prediction

We have experimented with different variants of our proposed task prediction method on the same task models and present the results on the CIFAR100/5 split in Table 6. We start off with the ensemble class prediction approach

Methods	Accuracy
Ensemble Class Prediction [12]	37.5
Entropy [44]	55.3
cross-entropy	54.9
∇ (cross-entropy) + mean filters	56.6
∇ (cross-entropy) + aug + mean filters	59.0
∇ (cross-entropy) + entropy aug	58.8
∇ (cross-entropy) + entropy aug + mean filters	59.4

Table 6. Average incremental accuracy till last task for different task prediction methods on CIFAR100/5 split.

of TCL [12] and show that it gives the lowest accuracy. Subsequently, we present the result for the vanilla entropy-based task prediction approach used in multiple previous works [44, 47]. Lastly, we explore different components of our gradient aggregation approach and show how the accuracy improves with every step. Our proposed approach leverages the gradients of cross-entropy, entropy weighted augmentations and mean filters leading to an absolute gain of 4.1% over the vanilla entropy baseline. It should be noted that entropy weighted augmentations significantly outperform vanilla augmentations for initial tasks. For example, entropy weighted augmentations lead to respective relative gains of 1.1% and 1.3% over vanilla augmentations for the initial $\lceil \frac{T}{2} \rceil$ tasks (minus first task which has same accuracy for both baselines) on CIFAR100/5 ($T=5$) and CIFAR100/10 ($T=10$) splits.

6. Conclusion

In this work, we propose a highly efficient expansion-based model that accumulates all the previously learned information. The expansion rate is dynamic and it depends on the task complexity. The proposed expansion approach uses a simple and efficient in-layer *filter and channel* expansion which is generic and can be applied to any convolutional network. We also propose a replay-free, flat minima based task prediction strategy which can be used to predict a large number of task sequences. The approach shows promising results for the TIL, CIL and generative continual learning settings. In expansion-based approaches, task prediction is the key limitation and it will be interesting to explore the same in future.

References

- [1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory Aware Synapses: Learning What (not) to Forget. *European Conference on Computer Vision*, 2018. 1, 2, 5, 6
- [2] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *ICLR*, 2020. 5
- [3] 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
- [4] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence. *European Conference on Computer Vision*, 2018. 1, 6
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A Large-Scale Hierarchical Image Database. *Computer Vision and Pattern Recognition*, 2009. 5
- [6] Qiang Gao, Zhipeng Luo, and Diego Klabjan. Efficient architecture search for continual learning. *arXiv preprint arXiv:2006.04027*, 2020. 2
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. *Neural information processing systems*, 2014. 7
- [8] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *Computer Vision and Pattern Recognition*, 2016. 1, 5
- [9] Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 9(1):1–42, 1997. 3
- [10] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, pages 2554–2558, 1982. 2
- [11] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, Picking and Growing for Unforgetting Continual Learning. *NeurIPS*, 2019. 2
- [12] Gyuhak Kim, Changnan Xiao, Tatsuya Konishi, Zixuan Ke, and Bing Liu. A theoretical study on solving continual learning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 2, 8
- [13] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming Catastrophic Forgetting in Neural Networks. *Proceedings of the National Academy of Sciences*, 2017. 1, 2, 5, 6, 7
- [14] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. 2009. 5, 7
- [15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 5
- [16] Abhishek Kumar, Sunabha Chatterjee, and Piyush Rai. Nonparametric Bayesian Structure Adaptation for Continual Learning. *arXiv preprint arXiv:1912.03624*, 2019. 2
- [17] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015. 5, 6
- [18] Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A Neural Dirichlet Process Mixture Model for Task-Free Continual Learning. *ICLR*, 2020. 2
- [19] Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. In *Advances in neural information processing systems*, pages 4652–4662, 2017. 2
- [20] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017. 2, 5, 6
- [21] Xiaolei Liu, Chenshen Wu, Mikel Menta, Luis Herranz, Bogdan Raducanu, Andrew D Bagdanov, Shangling Jui, and Joost van de Weijer. Generative feature replay for class-incremental learning. In *CVPR Workshops*, pages 226–227, 2020. 2
- [22] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–82, 2018. 2, 5
- [23] Arun Mallya, Svetlana Lazebnik, dummy dummy, and dummy dummy. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 2, 5
- [24] Marc Masana, Tinne Tuytelaars, de van, and Joost Weijer. Ternary feature masks: continual learning without any forgetting. *CVPR-Workshop*, 2021. 2, 5, 6
- [25] Nikhil Mehta, Kevin J Liang, Vinay Kumar Verma, and Lawrence Carin. Continual Learning using a Bayesian Nonparametric Dictionary of Weight Factors. *Artificial Intelligence and Statistics*, 2021. 2
- [26] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. *NeurIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011. 7
- [27] German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual Lifelong Learning with Neural Networks: A Review. *Neural Networks*, 2019. 2
- [28] G. Petit, A. Popescu, H. Schindler, D. Picard, and B. Delezoide. Fertil: Feature translation for exemplar-free class-incremental learning. *WACV*, 2023. 2, 5, 6
- [29] Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random Path Selection for Continual Learning. *Neural Information Processing Systems*, 2019. 2
- [30] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml: An incre-

- mental task-agnostic meta-learning approach. *arXiv preprint arXiv:2003.11652*, 2020. 1, 2
- [31] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. iCaRL: Incremental Classifier and Representation Learning. *Computer Vision and Pattern Recognition*, 2017. 1, 2
- [32] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience Replay for Continual Learning. *Neural Information Processing Systems*, 2019. 2
- [33] Anurag Roy, Vinay K Verma, Sravan Voonna, Kripabandhu Ghosh, Saptarshi Ghosh, and Abir Das. Exemplar-free continual transformer with convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5897–5907, 2023. 1, 2
- [34] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive Neural Networks. *arXiv preprint arXiv:1606.04671*, 2016. 1, 2, 5
- [35] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557, 2018. 2, 5
- [36] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual Learning with Deep Generative Replay. *Neural Information Processing Systems*, 2017. 2
- [37] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017. 1
- [38] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556*, 2014. 5
- [39] Pravendra Singh, Vinay Kumar Verma, Pratik Mazumder, Lawrence Carin, and Piyush Rai. Calibrating cnns for lifelong learning. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 5
- [40] Habib Slim, Eden Belouadah, Adrian Popescu, and Darian Onchis. Dataset knowledge transfer for class-incremental learning without memory. In *WACV, 2022*. 5, 6
- [41] Gido M van de Ven, Hava T Siegelmann, and Andreas S Tolias. Brain-inspired replay for continual learning with artificial neural networks. *Nature communications*, pages 1–14, 2020. 2
- [42] 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
- [43] Sakshi Varshney, Vinay Kumar Verma, PK Srijith, Lawrence Carin, and Piyush Rai. Cam-gan: Continual adaptation modules for generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:15175–15187, 2021. 7
- [44] Vinay Kumar Verma, Kevin J Liang, Nikhil Mehta, Piyush Rai, and Lawrence Carin. Efficient feature transformations for discriminative and generative continual learning. In *CVPR*, pages 13865–13875, 2021. 1, 2, 5, 6, 7, 8
- [45] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. In *International Conference on Learning Representations*, 2020. 2
- [46] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-UCSD Birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010. 5
- [47] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *Advances in Neural Information Processing Systems*, 33, 2020. 1, 2, 4, 7, 8
- [48] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory Replay GANs: Learning to Generate New Categories without forgetting. *Neural Information Processing Systems*, pages 5962–5972, 2018. 2, 5, 7
- [49] Ju Xu and Zhanxing Zhu. Reinforced continual learning. In *Advances in Neural Information Processing Systems*, pages 899–908, 2018. 2
- [50] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021. 1, 5
- [51] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Scalable and Order-robust Continual Learning with Additive Parameter Decomposition. *International Conference on Learning Representations*, abs/1902.09432, 2020. 1, 2, 5, 7
- [52] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong Learning with Dynamically Expandable Networks. *International Conference on Learning Representations*, 2018. 1, 2
- [53] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. LSUN: Construction of a Large-scale Image Dataset using Deep Learning with Humans in the Loop. *arXiv preprint arXiv:1506.03365*, 2015. 5
- [54] Lu Yu, Bartłomiej Twardowski, Xialei Liu, Luis Herranz, Kai Wang, Yongmei Cheng, Shangling Jui, and Joost van de Weijer. Semantic drift compensation for class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6982–6991, 2020. 5, 6
- [55] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning Through Synaptic Intelligence. *International Conference on Machine Learning*, 2017. 2, 5, 6
- [56] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. StackGAN++: Realistic Image Synthesis with Stacked Generative Adversarial Networks. *Transactions on Pattern Analysis and Machine Intelligence*, 2018. 5, 7
- [57] Junting Zhang, Jie Zhang, Shalini Ghosh, Dawei Li, Serafettin Tasci, Larry Heck, Heming Zhang, and C-C Jay Kuo. Class-Incremental Learning via Deep Model Consolidation. *Winter Conference on Applications of Computer Vision*, 2020. 5, 6
- [58] Jeffrey O Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: Network adaptation via additive side networks. *arXiv preprint arXiv:1912.13503*, 2019. 1, 2

- [59] Yang Zhao, Hao Zhang, and Xiuyuan Hu. Penalizing gradient norm for efficiently improving generalization in deep learning. 2022. [3](#)
- [60] Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, and De-Chuan Zhan. Pycil: A python toolbox for class-incremental learning, 2021. [5](#)
- [61] Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *NeurIPS*, 34:14306–14318, 2021. [2](#), [5](#), [6](#)
- [62] Kai Zhu, Wei Zhai, Yang Cao, Jiebo Luo, and Zheng-Jun Zha. Self-sustaining representation expansion for non-exemplar class-incremental learning. In *CVPR*, pages 9286–9295, 2022. [2](#), [5](#), [6](#)