

GATED CLASS-ATTENTION WITH CASCADED FEATURE DRIFT COMPENSATION FOR EXEMPLAR-FREE CONTINUAL LEARNING OF VISION TRANSFORMERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Vision transformers (ViTs) have achieved remarkable successes across a broad range of computer vision applications. As a consequence there has been increasing interest in extending continual learning theory and techniques to ViT architectures. In this paper, we propose a new method for exemplar-free class incremental training of ViTs. The main challenge of exemplar-free continual learning is maintaining plasticity of the learner without causing catastrophic forgetting of previously learned tasks. This is often achieved via exemplar replay which can help recalibrate previous task classifiers to the feature drift which occurs when learning new tasks. Exemplar replay, however, comes at the cost of retaining samples from previous tasks which for some applications may not be possible. To address the problem of continual ViT training, we first propose *gated class-attention* to minimize the drift in the final ViT transformer block. This mask-based gating is applied to class-attention mechanism of the last transformer block and strongly regulates the weights crucial for previous tasks. Secondly, we propose a new method of *feature drift compensation* that accommodates feature drift in the backbone when learning new tasks. The combination of gated class-attention and cascaded feature drift compensation allows for plasticity towards new tasks while limiting forgetting of previous ones. Extensive experiments performed on CIFAR-100 and Tiny-ImageNet demonstrate that our method outperforms existing exemplar-free state-of-the-art methods without the need to store any representative exemplars of past tasks.

1 INTRODUCTION

The initial excellent results of transformers for language tasks (Vaswani et al., 2017) have encouraged its application also for vision applications (Dosovitskiy et al., 2020). Vision Transformers (ViTs) currently achieve excellent results for many applications (Strudel et al., 2021; Liu et al., 2021; Caron et al., 2021). Most existing work on ViT training assumes that all training data is jointly available, an assumption which does not hold for real-world applications in which data arrives in a sequence of non-overlapping tasks. Continual learning considers learning from a non-IID stream of data. Applying a naive finetuning approach to such data results in a phenomenon called *catastrophic forgetting* which results in a drastic drop in performance on previous tasks (Goodfellow et al., 2014). The main goal of continual learning algorithms is to maximize the stability-plasticity trade-off (Mermillod et al., 2013), i.e. to mitigate forgetting of previously learned classes while maintaining the plasticity required to learn new ones.

One of the most successful approaches to preventing forgetting of previous tasks is *exemplar rehearsal* in which a subset of images from previous tasks is stored and then rehearsed when learning new ones (Rebuffi et al., 2017; Riemer et al., 2018; Chaudhry et al., 2018b; Lopez-Paz & Ranzato, 2017; Buzzega et al., 2020; 2021; Bang et al., 2021). Because of its success, the rehearsal technique has also been adopted by the initial works on continual learning for ViTs (Douillard et al., 2022; Wang et al., 2022a). However, for many applications the storage of previous task data might not be possible. This is especially true for applications with strict memory constraints and those where privacy or data use legislation prevents the long-term storage of data. In order to overcome this limitation, exemplar-free methods have been investigated (Kirkpatrick et al., 2017; Li & Hoiem, 2017;

Yu et al., 2020; Yan et al., 2021). These methods do not store any data from previous tasks, however their application to continual learning of ViTs has not been fully explored.

In this paper we propose one of the first exemplar-free methods for class-incremental learning (CIL). One of the challenges of exemplar-free continual learning is that these models tend to forget previously learned features while they are learning features for new tasks. The architecture we propose is based on a gated class-attention mechanism applied to the ViT decoder in order to mitigate the forgetting of learned features. As proposed by Serra et al. (2018), mask-based gating mechanisms are usually associated with the task-incremental scenario (i.e. scenarios in which a task id is known at inference time). We propose a solution to overcome this limitation by applying the masking mechanism only on the transformer decoder via multiple forward passes. This solution allows us to use mask-based gating in a class-incremental setup.

Mask-based gating prevents the drift of weights in the decoder, however it does not mitigate the drift in the transformer encoder. In fact, learning a stable backbone in the exemplar-free scenario is very difficult due to the drift in encoder weights that occurs during the learning of new tasks. To address this, we propose a method for backbone regularization in combination with a feature drift compensation mechanism that uses a cascade of projection networks that map the current backbone features to those of the previous backbone. This allows increased plasticity while maintaining stability across tasks, incurring a small computational cost due to the feature projection cascade. We also show, however, that knowledge distillation can be used to alleviate the computational burden of the projection cascade and the need for multiple forward passes at inference time.

The main contribution of this work are:

- a gated class-attention mechanism, called GCAB, that mitigates weight drift in the transformer decoder while also overcoming the need for a task id at inference time;
- a novel method for backbone regularization and feature drift compensation applied to the transformer encoder that increases plasticity towards new classes while maintaining stability in the previously learned ones;
- a method for GCAB distillation that reduces the computational overhead due to multiple forward passes and the memory overhead of storing projection networks for each task;
- experiments on multiple CIL benchmarks that demonstrate that our approach achieves state-of-the-art performance when compared to other exemplar-free methods.

2 METHOD

In this section, we propose our approach to exemplar-free ViT-based class-incremental learning. We begin by introducing the problem setup for class-incremental learning and the transformer architecture used in subsequent sections to perform exemplar-free class-incremental learning with Vision Transformers.

2.1 PROBLEM SETUP

Here we define the class-incremental learning setup and the specific Vision Transformer architecture we use.

Class-incremental learning setup. In class-incremental learning the model must learn a sequence of T tasks, where each task t introduces a number of new classes C^t . The data D^t of task t contains samples (x_i, y_i) , where x_i is input data labeled by $y_i \in C^t$. Note that we consider the case in which that there is no overlap between different task label sets: $C^i \cap C^j = \emptyset$ if $i \neq j$, as is commonly assumed (Masana et al., 2020). The model is evaluated on all previously seen classes $C^{\leq t} = \cup_{t' \leq t} C^{t'}$. Class-incremental learning differs from task-incremental learning in that it has no access to the task label t at inference time, and is therefore considered a more challenging setting (van de Ven & Tolias, 2018; Delange et al., 2021). Furthermore, in this paper, we consider the more restrictive setup of exemplar-free class-incremental learning in which no data from previous tasks is saved.

Transformer architecture. We use a vision transformer based on the one proposed by Dosovitskiy et al. (2020) and the recent improvements of Touvron et al. (2021). It consists of transformer encoder and decoder, each built with several multi-head attention blocks. In Figure 2 (left) we give a

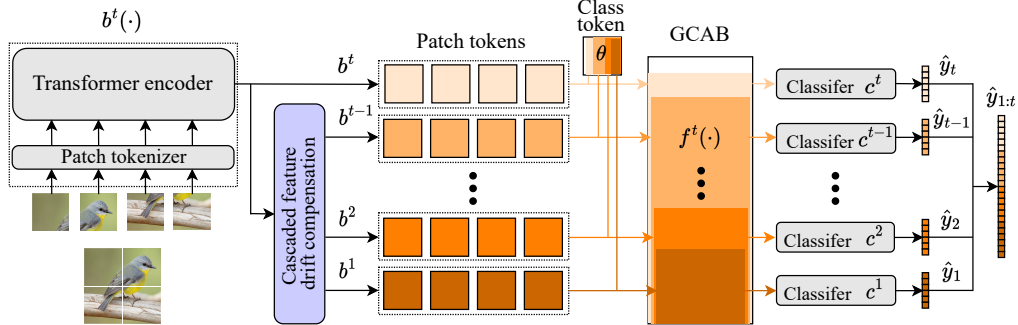


Figure 1: Overview of the architecture. Our main contributions are the *gated class-attention block (GCAB)* to prevent forgetting in the final ViT block (Section 2.2) and the *cascaded feature drift compensation* to compensate for feature drift of the backbone network (Section 2.5).

schematic diagram of our architecture. Formally, the input image $x \in \mathbb{R}^{H \times W \times C}$ is passed through a patch tokenizer that splits x into N patches and projects them using a 2D convolutional layer to obtain a set of N patch tokens $x_0 \in \mathbb{R}^{N \times D}$. A learnable position embedding is added to the patch tokens as in (Gehring et al., 2017). The patch tokens x_0 are passed as input to a sequence of M transformer encoder blocks, each yielding tensors of the same dimensions. Each block is composed of a multi-head self-attention (SA) mechanism (Vaswani et al., 2017), layer normalization and a Multi-layer Perceptron (MLP), each with residual connections:

$$\begin{aligned} x'_l &= x_l + \text{SA}(x_l) \\ x_{l+1} &= x'_l + \text{MLP}(x'_l) \end{aligned} \quad (1)$$

We follow the design of CaiT (Touvron et al., 2021), and only insert a class token in combination with a class-attention layer in the last block of the decoder. In our solution, the transformer decoder is composed of one single block. To distinguish the various parts of the transformer network, we define the image output prediction $\hat{y} = c(f(b(x; \Psi)))$, where the backbone features $b(x; \Psi) \in \mathbb{R}^{(N+1) \times D}$ parameterized by Ψ are the output of the self-attention blocks, and $f(b(x; \Psi))$ refers to the feature output of the decoder before classifier c .

2.2 GATED CLASS-ATTENTION

The main idea behind parameter isolation methods is to isolate a limited set of parameters after learning each task (Delange et al., 2021; Mallya & Lazebnik, 2018; Rusu et al., 2016). We build upon the work of *Hard Attention to the Task* (HAT) (Serra et al., 2018) in which attention masks are learned for each task. The masks operate on the activations of the network. Parameters of the network used by previous tasks can be exploited by new tasks, thereby allowing for forward transfer, but their update is restricted to prevent forgetting. The main strengths of this approach are the good forward transfer with little or no forgetting of previous tasks, together with the ability to automatically learn which neurons to dedicate to each task within the capacity limit of the neural network.

The forward pass of parameter isolation methods is usually conditioned on the task, and is therefore restricted to task-incremental learning. A possible way to extend these methods to class-incremental learning would be to run one forward pass for each task and then combine the task predictions (e.g., by simple concatenation). However, this would increase the run time linearly in the number of tasks. To limit computational overhead, we propose to only apply attention masks to the last block of the ViT. In Section 2.5, we investigate a distillation procedure to further reduce computational overhead.

Mask-based class-attention gating. We apply the attention-gating in the last transformer block, i.e. the decoder, which contains class-attention (Touvron et al., 2021). This block combines the patch tokens from previous blocks with a learnable class token θ . In Figure 2 (left) we give a schematic diagram of the gated class-attention block. We define a number of learnable masks m for task t as:

$$m^t = \sigma(sAt^T), \quad (2)$$

where, slightly abusing notation, t represents both a task *index* and a *one-hot vector* identifying the current task, A is an embedding matrix, σ is the sigmoid activation function, and s is a positive scaling parameter. Here m^t refers to the neurons that have been selected for task t .

We can now define the gated forward pass through the final class-attention block by introducing a mask for all the activations contributing to the final class token output. These masks correspond to: the input tokens (m_i^t), queries and keys (m_{QK}^t), values (m_V^t), the MLP (m_1 and m_2), and the class-attention output (m_o). We can then compute the query (Q), key (K), value (V), attention (A) and self attention output (O) given the block token inputs O^t and these masks:

$$\begin{aligned} p &= [\theta, b] \in \mathbb{R}^{(N+1) \times D} \\ Q^t &= W_q(\theta \odot m_i^t), \quad K^t = W_k(p \odot m_i^t), \quad V^t = W_v(p \odot m_i^t) \\ A^t &= \text{softmax} \left\{ \left((Q^t \odot m_{QK}^t) (K^t \odot m_{QK}^t)^T \right) / \sqrt{d/h} \right\} \\ O^t &= W_o A^t (V^t \odot m_V^t). \end{aligned} \quad (3)$$

The gating mechanism is also applied to the MLP (see Eq. 1) according to:

$$\begin{aligned} b^t &= O^t + \theta \\ u^t &= W_1(b^t \odot m_1^t), \quad v^t = W_2(u^t \odot m_2^t) \\ f^t &= v^t + O^t \end{aligned} \quad (4)$$

We call this the Gated Class-Attention Block (GCAB). The masks are learned during the training of task t and their role is twofold: they select those activations that are used to compute the task-conditioned output and they restrict the backpropagation of future tasks, preventing changes to weights that used by previous tasks. Given that both m_i^t and m_1^t operate on the class token embedding θ , we set $m_1^t = m_i^t$. In the experimental section, we show that good results are obtained when sharing the weights and setting $m_{QK} = m_V = m_1 = m_o = m_i$, resulting in only two learnable masks m_i and m_2 . In Figure 2 (left) we show the interaction between the masks m_i and m_2 and the class-attention block.

For each task a dedicated classifier c^t is added which produces predictions $\hat{y}^t = W_{clf}(f^t \odot m_o^t)$ using the vector f^t output from the GCAB. We perform multiple forward passes of patches extracted from the backbone b^t through the decoder f^t (i.e. we pass it t times). For each pass s , the obtained vectors l^s are then passed to the corresponding classifier c^s . The classifier outputs are then concatenated $C = [c^1, \dots, c^t]$ and the binary cross entropy loss \mathcal{L}_{BCE} is computed using the target vector.

Training. During the training of the current task t the scaling parameter s from equation 2, is scaled with the batch index: $s = \frac{1}{s_{max}} + (s_{max} - \frac{1}{s_{max}}) \frac{i-1}{I-1}$, where i current batch index and I is the total number of batches in an epoch. This was found to be beneficial in (Serra et al., 2018).

At the end of the current task, the learned masks are accumulated as $m_*^{<t} = \max(m_*^{t-1}, m_*^{<t-1})$, where ‘*’ stands for any of the specific mask subscripts introduced above. Accumulated masks $m_*^{<t}$ are used during backpropagation to prevent updating weights considered important for the tasks observed so far. The masks are learned by minimizing the following loss function:

$$\mathcal{L}_{GCAB} = \lambda_{GCAB} \frac{\sum_x m_x^t (1 - m_x^{<t})}{\sum_x 1 - m_x^{<t}}, \quad (5)$$

where m_x^t is the mask learned at the current task t for component x of the GCAB, x ranges over the mask subscripts described above, and $m_x^{<t}$ is the cumulative mask. λ_{GCAB} is a tunable hyperparameter controlling the capacity of the masks learned during the tasks. This equation encourages new task mask m_x^t to be sparse, however it permits use of activations already used by previous tasks $m_x^{<t}$ at no cost.

The cumulative masks also play a pivotal role during the training of new tasks. Consider, for example, weights W_q that map from the input tokens (masked by m_i^t) to the queries (masked by m_{QK}^t). We then define the elements of the weight mask according to $M_{q,kl}^{<t} = 1 - \min(m_{i,k}^{<t}, m_{QK,l}^{<t})$ where $m_{i,k}^{<t}$ refers to the k -th element of $m_i^{<t}$. The update rule for the backpropagation of the gradient is then: $W_q = W_q - \lambda M_q^{<t} \odot \frac{\partial L}{\partial W_q}$. This update rule prevents the updating of part of the weights learned for previous tasks. The input mask also influences the updating of the class token embedding which is given by $W_\theta = W_\theta - \lambda(1 - m_i^{<t}) \odot \frac{\partial L}{\partial W_\theta}$. In Appendix A.2 we show the update rules for all weight matrices in the class-attention block.




Figure 2: (left) GCAB applies masks to the activations of the class-attention block. This allows prevention of forgetting of previous tasks. (right) The cascade of feature projection networks maps the current backbone features to previous task features to compensate for feature drift.

2.3 BACKBONE REGULARIZATION AND CASCADED FEATURE DRIFT COMPENSATION

In the previous section we applied gating only to the last transformer block to limit computational overhead. Gating ensures that only minimal changes occur to the weights relevant to previous tasks, however the network can still suffer from forgetting due to backbone feature drift. To mitigate forgetting in the backbone, we apply regularization to these features.

A straightforward way to prevent forgetting in the backbone network via regularization is feature distillation (Hou et al., 2019). Feature distillation encourages backbone features at task t to remain close to those at task $t-1$, however this was found to limit plasticity (Douillard et al., 2020). To ensure stability without sacrificing plasticity, some recent works in continual learning of self-supervised representations have proposed to learn a projector between feature extractors (Fini et al., 2022; Gomez-Villa et al., 2022). This approach, called Projected Functional Regularization (PFR), introduces a projection network p^t that maps the current backbone features to those of the previous backbone. PFR allows the new backbone to learn new features without imposing a high regularization penalty as long as the new features can still be projected back to those of the previous backbone. The PFR loss function is:

$$\mathcal{L}_{\text{pfr}} = \lambda_{\text{pfr}} \mathbb{E}_{x \sim \mathcal{D}^t} [\mathcal{S}(p^t(b(x; \Psi^t)), b(x; \Psi^{t-1}))] \quad (6)$$

where \mathcal{S} is the cosine distance, λ_{pfr} is a trade-off parameter, and Ψ^t refers to the parameters of the backbone after learning task t .

The gained plasticity induced by PFR leads to a misalignment of the current backbone with previous class-attention layers and classifiers. Consider the predictions of the system after training t tasks on data of task $s \leq t$: $\hat{y} = c^t(f^t(b^t(x)))$. The gating mechanism described above ensures that $c^t(f^t(\cdot)) \approx c^s(f^s(\cdot))$, however $b^t \neq b^s$. That is, the backbone feature representation at task t has *drifted* away from the representation at previous task s . To address this we perform *feature drift compensation* (see Fig. 2 (right)) and exploit the learned projection networks p^t in a *cascade* to align the current backbone with the learned class-attention block at any previous task s :

$$\hat{y}_s = c^t(f^t(p^{s+1}(p^{\dots}(p^{t-1}(p^t(b^t(x))))))) \quad (7)$$

All projection networks p^t must be saved in this formulation in order to compute the projection cascade that compensates for backbone feature drift. In Section 2.5 we show how knowledge distillation can be used to eliminate the need for projection networks p^t at inference time,

2.4 TRAINING OBJECTIVE AND INFERENCE

The final objective we use for incremental training of the model is the sum of the three loss functions:

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \mathcal{L}_{\text{pfr}} + \mathcal{L}_{\text{GCAB}} \quad (8)$$

After training the current task, an evaluation phase is performed over all classes seen so far. At task t , the model is tested for the tasks $t' \leq t$. All images are passed through the backbone to obtain tokens b . The tokens are passed t times through the Gated Class-Attention Block. Based on the task index of the forward pass t' , the composition of the previously stored projection networks is used as in Equation 7 to align the current backbone features to those of previous task t' . In Figure 2 we give a schematic diagram of the Feature Drift Compensation mechanism during inference. During inference the parameter s of Equation 2 is equal to s_{max} .

2.5 GCAB DISTILLATION

To overcome the increased computational cost due to multiple forward passes for all tasks and the cascaded projection layers, we perform knowledge distillation (Hinton et al., 2015) to transfer the class-conditioned GCAB (the *teacher*) into a single class-attention block (CAB) with the same architecture as the GCAB but *without* masks, and an aggregated classifier c^{A_t} (*student*). This reduces the number of parameters, since the task projection networks are no longer needed, and eliminates the need for multiple forward passes. As shown in Fig. 3, the CAB and c^{A_t} trained by minimizing the Kullback–Leibler (KL) divergence between the logits output by the teacher and student models. Note that this distillation is conducted only with the data from task t , while the transformer backbone b and teacher hyper-classifier are frozen during training. We also investigate the use of static masks in the CAB to leave unused capacity in the student network in order to accommodate potential future tasks (see Section 4.4 for details).

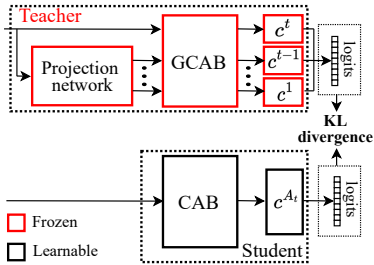


Figure 3: GCAB distillation.

3 RELATED WORK

Continual Learning. Continual learning algorithms can be grouped in three categories (Masana et al., 2020): regularization approaches, parameter based regularization (Kirkpatrick et al., 2017; Liu et al., 2018; Lee et al., 2020; Zenke et al., 2017), and data based regularization (Li & Hoiem, 2017; Hou et al., 2019; Wu et al., 2019; Castro et al., 2018; Jung et al., 2016; Dhar et al., 2019; Zhang et al., 2020); rehearsal approaches, which store (Rebuffi et al., 2017; Chaudhry et al., 2018a) or generate exemplars (Wang et al., 2021; Zhai et al., 2021); and bias-correction approaches (Castro et al., 2018; Wu et al., 2019; Hou et al., 2019).

Continual Learning with ViTs. Visual Transformers recently outperforms convolutional neural networks and in particular resnet(He et al., 2016) in several tasks like classification (Dosovitskiy et al., 2020) or segmentation (Zheng et al., 2021). Although ViTs are considered state-of-the-art models, their application in continual learning has not been fully explored. Douillard et al. (2022) proposed a transformer-based architecture, called DyTox, with a dynamic task-token expansion for mitigating catastrophic forgetting. Wang et al. (2022a) proposed an inter-task attention mechanism for ViTs. Wang et al. (2022b) described a prompting method for continually learning a classifier using a pretrained, frozen ViT backbone. Even if the performances showed in these work are very remarkable, the challenge of continually learning the parameters of a ViT without storing exemplars or using a pretrained model, is still open. Differently from previous work, we propose an exemplar-free ViT approach to class-incremental learning. Our work is inspired by DyToX (Douillard et al., 2022), which applies a task conditioned class-attention block. However, DyToX shares the class-attention block parameters between tasks which leads to forgetting, and it therefore requires exemplars to counter this. Instead, we replace the task token with a task-specific gating function that prevents forgetting and does not require exemplars. In addition, we introduce feature drift compensation, which allows for more plasticity in the backbone.

Parameter Isolation in Continual Learning. In this family of algorithms, learnable masks are applied to the weights of the model in order to reduce forgetting. Mallya et al. (2018) proposed Piggyback, a masked-based method able to learn the weight masks while training a backbone. The same group proposed Packnet (Mallya & Lazebnik, 2018) which, via iterative pruning and sequential re-training, is able to add multiple tasks to a single network. Serra et al. (2018) proposed to apply masks to layer activations in order to limit the update of the parameters more relevant to a specific task. Masana et al. (2021) proposed a system of ternary-masks applied on to layer activations for preventing catastrophic forgetting and backward transfer. Yan et al. (2021) proposed a Dynamical Expandable Representation (DER) for continual learning. In this work, channel-level masks are used for pruning the feature extractor. Rajasegaran et al. (2019) proposed Random Path Selection (RPS). This approach uses a parameter isolation mechanism, distillation, and a replay-buffer to learn different paths for each task without the need for a task id during inference.

Exemplar-Free Continual Learning. This is one of the most challenging scenario in continual learning. In this paradigm, it is not possible to store any exemplars from the previously observed

Method	# Params	Exemplar-Free	CIFAR-100 5 Tasks		CIFAR-100 10 Tasks		Tiny-ImageNet	
			Class-IL	Task-IL	Class-IL	Task-IL	Class-IL	Task-IL
ER (Riemer et al., 2018)	11.2	✗	21.94	62.41	14.23	67.57	8.79	39.16
AGEM (Chaudhry et al., 2018b)	11.2	✗	17.97	53.55	9.44	55.04	8.28	23.79
iCaRL (Rebuffi et al., 2017)	11.2	✗	30.12	55.70	22.38	60.81	8.64	28.41
FDR (Benjamin et al., 2018)	11.2	✗	22.84	63.57	14.85	65.88	8.77	40.15
DER++ (Buzzega et al., 2020)	11.2	✗	27.46	62.55	21.76	59.54	11.16	40.91
ERT (Buzzega et al., 2021)	11.2	✗	21.61	54.75	12.91	58.49	10.85	39.54
RM (Bang et al., 2021)	11.2	✗	32.23	62.05	22.71	66.28	13.58	41.96
LVT (Wang et al., 2022a)	8.9	✗	39.68	66.92	35.41	72.80	17.34	46.15
DyToX (Douillard et al., 2022)	10.7	✗	49.10	-	34.83	-	13.14	-
Ours	12.4	✓	49.36	81.01	35.57	82.08	26.82	65.92
Ours (dis-80)	10.7	✓	48.85	79.75	35.42	81.97	26.44	65.02

Table 1: Comparison on CIFAR-100 and Tiny-ImageNet. All methods except ours use a memory buffer of 200 exemplars. The accuracies reported here are the ACC_{TAG} and ACC_{TAW} computed after the training the last task. *Ours (dis-80)* indicates using GCAB distillation with 80% capacity usage in the distilled CAB.

classes. Li & Hoiem (2017) proposed an exemplar-free data regularization approach to mitigate forgetting. This method distills knowledge of the previous model into the new one in order to prevent weight drift while learning the new task data. Kirkpatrick et al. (2017) described an exemplar-free weight regularization approach called Elastic Weight Consolidation (EWC) for preventing weight drift. Similarly, Aljundi et al. (2018) proposed a method that accumulates the importance of each model parameter by analyzing the effect of their change to the predicted output. Yu et al. (2020) proposed a semantic drift compensation mechanism to compensate for feature drift in previous tasks by approximating it with the drift estimated with current task data. Toldo & Ozay (2022) presented a framework for modeling the semantic drift of model weights of and estimating feature drift in the representation of previously learned classes. Pelosin et al. (2022) proposed an attention distillation mechanism for exemplar-free visual transformer in task-incremental learning.

4 EXPERIMENTAL RESULTS

For our experiments we consider two datasets: CIFAR-100 (Krizhevsky et al., 2009) and Tiny-ImageNet (Le & Yang, 2015) (for details see Section A.1). We set the number of transformer encoder blocks to $M = 5$, each one with $H = 12$ heads for the multi-head self-attention mechanism. The dimension of the embeddings is set to $D = 384$. We train on each task for 500 epochs using Adam with $lr = 1e^{-4}$ and a batch size of 128. We set hyperparameters $\lambda_{pfr} = 0.001$, $\lambda_{GCAB} = 0.05$ and $s_{max} = 800$.

4.1 COMPARISON WITH THE STATE-OF-THE-ART

Equal task split scenarios. We consider two different CIL scenarios: 5 tasks and 10 tasks equally split among all classes. For CIFAR-100, tasks contain 20 classes for the 5 task scenario and 10 for the 10 task scenario. For the Tiny-ImageNet we consider only the 10-task scenario, with 20 classes in each task. For CIFAR-100 and Tiny-ImageNet the images are split in $N = 64$ patches. We use the task-agnostic top-1 accuracy over all the classes of the dataset after training the last task: $ACC_{TAG} = \frac{1}{N} \sum_{t=1}^N a_t$, where N is the total number of classes in the dataset. For the task-aware scenario we report the mean accuracy over all the tasks after training the last task: $ACC_{TAW} = \frac{1}{T} \sum_{t=1}^T a_t$, where T is the total number of tasks. We compare our approach with several continual learning methods: ER (Riemer et al., 2018), AGEM(Chaudhry et al., 2018b), iCaRL(Rebuffi et al., 2017), FDR(Benjamin et al., 2018), DER++ (Buzzega et al., 2020), ERT(Buzzega et al., 2021), RM(Bang et al., 2021), LVT(Wang et al., 2022a) and DyToX (Douillard et al., 2022). The memory buffer for these methods is limited to 200 images, which is the setting proposed in (Wang et al., 2022a)).

From Table 1, we see that our method outperforms all methods in both class-incremental and task-incremental scenarios. Note that our approach outperforms the two other ViT-based methods, Dytox and LVT. For the Tiny-ImageNet, we observed that our method was able of outperforming all the compared methods by a large margin. Especially, we obtain a significant improvement of almost 10% with respect to LVT, while LVT uses exemplars and we do not.

Method	CIFAR-100 5 splits	Tiny-ImageNet 5 splits
EWC (Kirkpatrick et al., 2017)	26.26	14.63
LWF (Li & Hoiem, 2017)	39.51	40.62
LWM (Dhar et al., 2019)	40.49	28.39
PASS (Zhu et al., 2021)	56.53	47.00
SDC (Yu et al., 2020)	57.62	47.89
Evanescent (Toldo & Ozay, 2022)	59.37	48.56
Ours	60.85	46.35

Table 2: Comparison of exemplar-free methods on CIFAR-100 and Tiny-ImageNet on the larger first task scenario. The accuracies reported here are the average incremental accuracies ACC_{AVG} .

Gated Class Attention	Backbone Regularization	Feature Drift Compensation	Avg Acc
			28.30
✓			51.56
✓	FD		50.93
✓	PFR (2)		25.66
✓	PFR (1)	✓	52.98
✓	PFR (2)	✓	55.30

Table 3: Ablation study on the components of our architecture. Results are on the 10 task scenario on CIFAR-100. We report the average incremental accuracy ACC_{AVG} . In parentheses are the number of layers of the MLP used for PFR.

Larger first task scenarios. In this setting, the number of classes involved in the training of the architecture during the first task is 40 for CIFAR-100 and 100 for TinyImagenet. The remaining classes (60 and 100 respectively) are then split in 5 task respectively. This scenario is often used in exemplar-free CIL because it allows to train a high-quality backbone on the first tasks which does not need much training on the subsequent tasks.

We compare our model with several methods from the state-of-the-art. For this scenario, we considered as metric the average incremental accuracy defined as: $ACC_{AVG} = \frac{1}{T} \sum_{t=1}^T a_t$, where a_t is the task-agnostic accuracy computed over the classes observed up to task t (as used by (Toldo & Ozay, 2022) for this experiment). In Table 4.1, we observe that our method obtains the best performance on the CIFAR-100 dataset. For the TinyImagenet instead, we obtained a competitive performance compared to different method as (Yu et al., 2020; Toldo & Ozay, 2022; Zhu et al., 2021)¹.

4.2 ABLATION STUDY

We ablate on the importance of the different components of our approach. In Table 3 we report 6 possible configurations on the 10-task split of CIFAR-100. First, we consider fine-tuning our architecture without applying any continual learning strategy to the base architecture (28.3%). Then we apply the gated class-attention mechanism to the transformer decoder. This solution increases the average accuracy by 23%, showing the importance of preventing forgetting in the final block. As explained in Section 2.2, we use only two masks for gating the transformer decoder. This does not prevent backbone weight drift when passing from one task to the next. Therefore, we apply a simple backbone regularization based on feature distillation (Hou et al., 2019) which does not increase overall performance (50.93%). When we replace feature distillation with the projected function regularization in combination with feature drift compensation, we obtain better performance – notably, a more than 4% increase when using a 2-layer MLP. These results confirm the importance of projecting the learned backbone features to the previous features space.

We additionally analyze the design of the masks considered in the Gated Class-Attention Block. We consider the larger first task scenario, 40 classes in the first task, and the remaining are split equally split in 5 tasks. As reported in Figure 4, we plot the results obtained with two different choices for the masks: 1) For the orange curve $m_{QK} \neq m_V \neq m_1 \neq m_o \neq m_i \neq m_2$; 2) For the blue curve $m_{QK} = m_V = m_1 = m_o = m_i$. From the figure we see that using the two masks for the gated class-attention block provides the best results.

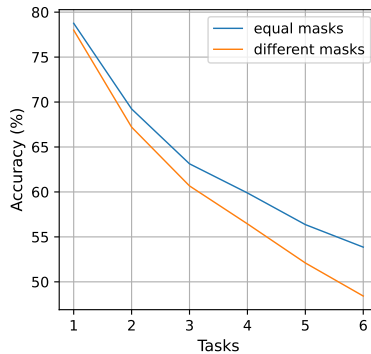


Figure 4: Ablation on the CIFAR-100 dataset in the larger first task scenario. The two curves show different choices for the Gated Class-Attention Block.

¹We hypothesize that the \mathcal{L}_{GCAB} is sensitive to non-uniform splitting of data, and that tailored hyperparameters would improve the results of our method (as in Table 1). We keep this, however, for future research.




Figure 5: t-SNE visualization after learning the first task (output of $f(\cdot)$) after the first task (stage 1) and second task (stage 2). In stage 2, we show the models learned without the Projection network (PN), with only PN during training, and with PN during both training and inference.




Figure 6: Average accuracy after the last task with and without GCAB distillation.

4.3 EMBEDDING VISUALIZATION

To verify the effectiveness of Cascaded Feature Drift Compensation, we analyzed the embeddings produced by the GCAB. In Figure 5 we show the the embedding produced by the GCAB in a two dimension using the t-SNE (Van der Maaten & Hinton, 2008). These visualizations are based on the task-agnostic embeddings in the 10 task scenario on CIFAR-100. The embedding of the images from the test set of the first task are shown after the training of the second task. These visualizations show that after training the second task, the cascaded projection network preserves clusters of first task classes.

4.4 DISTILLATION RESULTS

We conduct the experiments with GCAB distillation in four scenarios: CIFAR-100 with 5 and 10 tasks, and Tiny-ImageNet with 5 and 10 tasks. We freeze the transformer encoder, projection network, GCAB, and classifiers after the last task and train the student CAB (as shown in Figure 3) with an Adam optimizer and a learning rate $5e^{-3}$ for 200 epochs with only data from the last task. When performing GCAB distillation we use static binary masks at the same position as the masks in GCAB (see Fig.2 (right)) to control the capacity usage in the student CAB (this potentially allows us to continue training on further tasks). From the results in Figure 6 we see that GCAB distillation obtains almost the same performance as the model before distillation when the capacity usage is higher than 80%, and only a small performance drop at 60% capacity usage. As shown in Table 4, GCAB distillation can overcome the increased computational cost from multiple forward passes for all tasks during inference.

	Time (ms)	GFLOPs
GCAB	~ 30	1.648
Distilled GCAB	~ 21	0.606

Table 4: Run time and computational cost with and without GCAB distillation after the last task on the CIFAR-100 10 task scenario.

5 CONCLUSION

In this paper we presented a new exemplar-free approach to class-incremental Visual Transformer training. Our method, through the Gated Class-Attention Mechanism, achieves low forgetting by learning and masking important neurons for each task. High plasticity is ensured via the Backbone Regularization and Feature Drift Compensation using a cascade of feature projection networks. Through classifier distillation, we are able to overcome the limitations due to the multiple forward passes and computational overhead required by multiple forward passes through a task-conditioned network. Ours is one of the first effective approaches to exemplar-free, class-incremental training of ViTs.

REPRODUCIBILITY STATEMENT

The code of our method, written in PyTorch 1.11.0, will be released in case of acceptance of the paper. All the datasets used in this work are publicly available online

REFERENCES

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 139–154, 2018.
- Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8218–8227, 2021.
- Ari S Benjamin, David Rolnick, and Konrad Kording. Measuring and regularizing networks in function space. *arXiv preprint arXiv:1805.08289*, 2018.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, and Simone Calderara. Rethinking experience replay: a bag of tricks for continual learning. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 2180–2187. IEEE, 2021.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9650–9660, 2021.
- 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)*, pp. 233–248, 2018.
- Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 532–547, 2018a.
- Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*, 2018b.
- Matthias Delange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Greg Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5138–5146, 2019.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pp. 86–102. Springer, 2020.
- Arthur Douillard, Alexandre Ramé, Guillaume Couairon, and Matthieu Cord. Dytox: Transformers for continual learning with dynamic token expansion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9285–9295, 2022.

- Enrico Fini, Victor G Turrisi da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-supervised models are continual learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9621–9630, 2022.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *International conference on machine learning*, pp. 1243–1252. PMLR, 2017.
- Alex Gomez-Villa, Bartłomiej Twardowski, Lu Yu, Andrew D Bagdanov, and Joost van de Weijer. Continually learning self-supervised representations with projected functional regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3867–3877, 2022.
- Ian J. Goodfellow, Mehdi Mirza, Xia Da, Aaron C. Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6211>.
- 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*, pp. 770–778, 2016.
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
- Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 831–839, 2019.
- Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetting learning in deep neural networks. *arXiv preprint arXiv:1607.00122*, 2016.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- Janghyeon Lee, Hyeong Gwon Hong, Donggyu Joo, and Junmo Kim. Continual learning with extended kronecker-factored approximate curvature. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9001–9010, 2020.
- Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pp. 2262–2268. IEEE, 2018.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10012–10022, 2021.
- David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017.
- Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2018.

- 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)*, pp. 67–82, 2018.
- Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint arXiv:2010.15277*, 2020.
- Marc Masana, Tinne Tuytelaars, and Joost Van de Weijer. Ternary feature masks: zero-forgetting for task-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3570–3579, 2021.
- Martial Mermillod, Aurélia Bugajska, and Patrick Bonin. The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects, 2013.
- Francesco Pelosin, Saurav Jha, Andrea Torsello, Bogdan Raducanu, and Joost van de Weijer. Towards exemplar-free continual learning in vision transformers: an account of attention, functional and weight regularization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3820–3829, 2022.
- Jathushan Rajasegaran, Munawar Hayat, Salman H Khan, Fahad Shahbaz Khan, and Ling Shao. Random path selection for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 2001–2010, 2017.
- Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauero. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv*, 2016.
- Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pp. 4548–4557. PMLR, 2018.
- Albin Soutif-Cormerais, Marc Masana, Joost Van de Weijer, and Bartłomiej Twardowski. On the importance of cross-task features for class-incremental learning. *CoRR*, 2021.
- Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7262–7272, 2021.
- Marco Toldo and Mete Ozay. Bring evanescent representations to life in lifelong class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16732–16741, 2022.
- Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 32–42, 2021.
- Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. In *NeurIPS Continual Learning Workshop*, 2018.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

- Liyuan Wang, Kuo Yang, Chongxuan Li, Lanqing Hong, Zhenguo Li, and Jun Zhu. Ordisco: Effective and efficient usage of incremental unlabeled data for semi-supervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5383–5392, 2021.
- Zhen Wang, Liu Liu, Yiqun Duan, Yajing Kong, and Dacheng Tao. Continual learning with lifelong vision transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 171–181, 2022a.
- Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 139–149, 2022b.
- 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*, pp. 374–382, 2019.
- Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3014–3023, 2021.
- 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*, pp. 6982–6991, 2020.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pp. 3987–3995. PMLR, 2017.
- Mengyao Zhai, Lei Chen, and Greg Mori. Hyper-lifelonggan: scalable lifelong learning for image conditioned generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2246–2255, 2021.
- 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. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1131–1140, 2020.
- Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6881–6890, 2021.
- Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5871–5880, 2021.

A APPENDICES

A.1 DATASETS

The CIFAR-100 dataset is composed of 60,000 images, each 32×32 pixels and divided into 100 classes. Each class has 500 training and 100 test images. The Tiny-ImageNet is a reduced version of the original ImageNet dataset with 200 classes. The classes are split into 500 for training, 50 for validation, and 50 for test (for a total number of 120,000 images). The images are 64×64 pixels.

A.2 UPDATE RULES FOR GATED CLASS-ATTENTION

The backpropagation update rules are the following:

$$\begin{aligned} W_Q &= W_Q - \lambda M_q^{<t} \odot \frac{\partial L}{\partial W_q} & W_k &= W_k - \lambda M_k^{<t} \odot \frac{\partial L}{\partial W_k} \\ W_v &= W_v - \lambda M_v^{<t} \odot \frac{\partial L}{\partial W_v} & W_\theta &= W_\theta - \lambda(1 - m_{i,k}^{<t}) \odot \frac{\partial L}{\partial W_\theta} \\ W_o &= W_o - \lambda M_o^{<t} \odot \frac{\partial L}{\partial W_o} & W_1 &= W_1 - \lambda M_1^{<t} \odot \frac{\partial L}{\partial W_1} \\ W_2 &= W_2 - \lambda M_2^{<t} \odot \frac{\partial L}{\partial W_2} & W_{\text{clf}} &= W_{\text{clf}} - \lambda(1 - m_o^{<t}) \odot \frac{\partial \text{clf}}{\partial W_{\text{clf}}} \end{aligned}$$

and the weight masks are defined as:

$$\begin{aligned} M_{q,kl}^{<t} &= 1 - \min(m_{i,k}^{<t}, m_{QK,l}^{<t}) & M_{k,kl}^{<t} &= 1 - \min(m_{i,k}^{<t}, m_{QK,l}^{<t}) \\ M_{v,kl}^{<t} &= 1 - \min(m_{i,k}^{<t}, m_{V,l}^{<t}) & M_{o,kl}^{<t} &= 1 - \min(m_{i,k}^{<t}, m_{1,l}^{<t}) \\ M_{1,kl}^{<t} &= 1 - \min(m_{1,i}^{<t}, m_{2,l}^{<t}) & M_{2,kl}^{<t} &= 1 - \min(m_{2,i}^{<t}, m_{o,l}^{<t}) \end{aligned}$$

A.3 ADDITIONAL RESULTS

Accuracy Matrices. In Figures 7 and 8, we show the task aware and task agnostic accuracy matrices for the 5-task scenario on CIFAR-100. We observe that, through our gated class-attention block, we achieve stable accuracy on previous tasks in the task-aware scenario. In Figure 8 the task-agnostic matrix shows a drop in performance when learning new tasks. However, this behaviour is expected since the problem becomes increasingly complex when adding new tasks.




Figure 7: Task-Aware on CIFAR-100, 5 tasks. Figure 8: Task-Agnostic on CIFAR-100, 5 tasks.

Cumulative Metrics. To verify the accuracy and forgetting of our method, we compute the Cumulative Accuracy and Cumulative Forgetting as defined by Soutif-Cormerais et al. (2021). We see from Figure 9 that the cumulative accuracy of our model in the 10-task scenario on CIFAR-100 is stable without any drop when increasing the number of observed tasks. We report in Figure 10 the cumulative forgetting for each task. Even if there is a slight increase for early tasks, the cumulative forgetting does not increase above 4%.





Figure 9: Cumulative Accuracy on the 10-task scenario of CIFAR-100. Figure 10: Cumulative Forgetting on the 10-task scenario of CIFAR-100.

Gating Capacity. In Figure 11, we show the percentage of used masked capacity as tasks are added. The experiments were performed on CIFAR-100 dataset for 5- and 10-task scenarios. We




Figure 11: Gated Class-Attention Block capacity on CIFAR-100 in the 5- and 10-task scenarios.

observe that most most of the available capacity is used during the firsts tasks. For subsequent ones the percentage of occupied capacity is significantly lower compared to the first ones. From the 5-task scenario curve we that there is almost 10% of the capacity available for possible new incoming tasks. For the 10-task scenario, however, after completing 60% of the tasks the capacity is almost full and the last tasks are using less than 2% of capacity each.

A.4 ADDITIONAL EMBEDDING VISUALIZATIONS

We report here the complete visualization of the embeddings produced with and without the projection network cascade. In the first row of the Figure 12 we show the embedding obtained without using the projection cascade. In the second and third rows we show the embeddings obtained using the projection network cascade during training. The difference between the middle and bottom row is the use of the Feature Drift Compensation mechanism during inference. The different columns of the figure show the embeddings produced by the network during the different inference phases. As previously explained, the projection network is introduced only during the training of the second

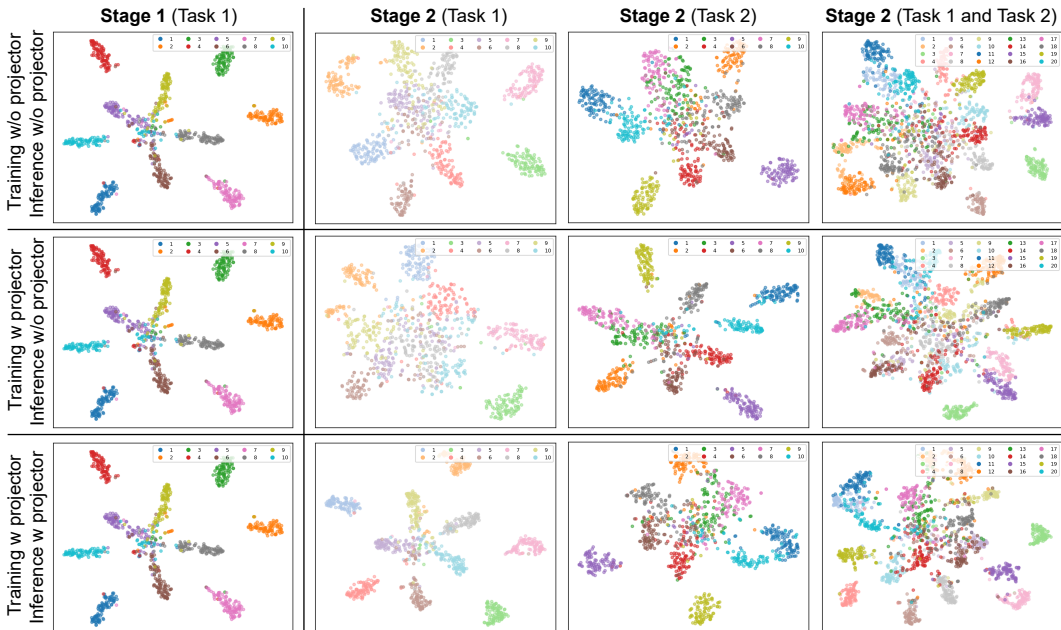


Figure 12: T-SNE visualization of embedding space (output of $f(\cdot)$) at the first two tasks (task-agnostic). **Top Row**: the model is learned without (w/o) projection network. **Middle Row**: projection network is used during training but not during inference. **Bottom Row**: projection network is involved in both training and inference.

task. For this reason Stage 1 of the three different presented strategies have the same embedding. The projection network learned during the training of the second task is used to align the embeddings produced by the current backbone to the ones produced by the previous one. We see that during inference on first task data after the training of the second task (i.e. Stage 2 in Figure 12) the Feature Drift Compensation better preserves first-task clusters in embedding space: using the projection network, the embeddings of first task classes are clearly clustered. We also analyze the test set embedding features of the second task (third column), after the training of the second task. In Figure 13 the task-aware feature embeddings with and without the use of the Feature Drift Compensation during the inference are reported.

A.5 COMPARISON OF OUR CONTRIBUTION WITH OTHER METHODS.

To clarify our contribution in this paper, we summarize the difference compared with recent continual learning methods in Table. 5 and Table. 6. As shown in Table. 5, our method is the first exemplar-free class incremental learning method based on the transformer. Pelosin et al. (2022) also propose a transformer-based CL method with attention distillation, but for task-incremental learning. The method of Wang et al. (2022b) requires a pre-trained backbone and dynamically prompts it to a sequence of tasks. Douillard et al. (2022) propose a transformer-based method with dynamic token expansion for continual learning. This method achieves very good performances on CIL, but requires exemplars for mitigating the forgetting. Finally (Wang et al., 2022a) propose a lifelong vision transformer without any pretraining but requiring exemplars.

In Table 6 we report recent parameter-isolation continual learning methods. To the best of our knowledge, ours is the first exemplar-free class-incremental learning method trained completely from scratch. Piggyback (Mallya et al., 2018) uses task-specific binary masking for adapting the backbone network to new tasks. Although it does not use exemplar rehearsal, it is based on a pretrained backbone. Similarly, Mallya & Lazechnik (2018) describe PackNet which is based on task-aware iterative pruning and re-training to mitigate forgetting and increase plasticity. These methods apply the learned masks directly to the weights of the network. Another category of parameter-isolation methods instead apply the masks to the layer activations. Serra et al. (2018) describe a task-

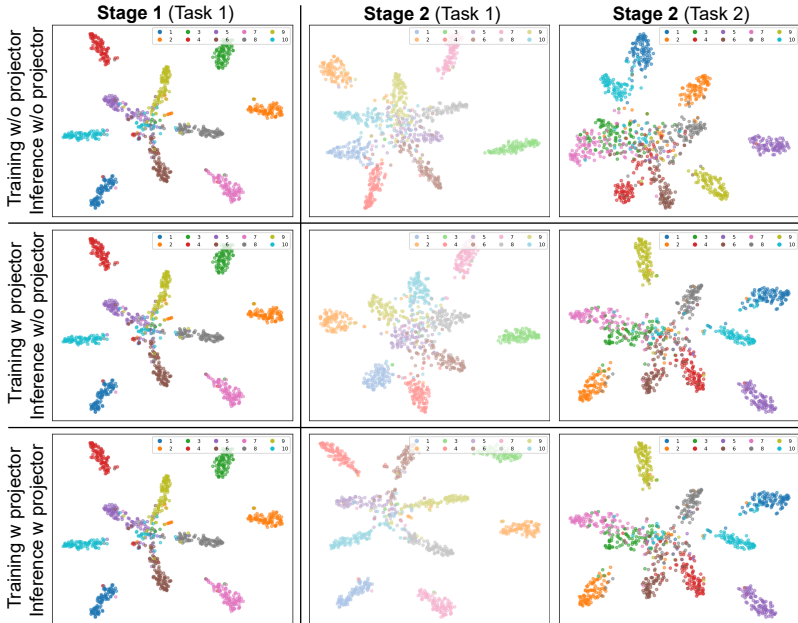


Figure 13: T-SNE visualization of embedding space (output of $f(\cdot)$) at the first two tasks (task-aware). **Top Row**: the model is learned without (w/o) projection network. **Middle Row**: projection network is used during training but not during inference. **Bottom Row**: projection network is involved in both training and inference.

Method	Architecture	Exemplars-Free	Non-Pretrained
Ours	ViT	✓	✓
Dytox (Douillard et al., 2022)	ViT	✗	✓
LVT (Wang et al., 2022a)	ViT	✗	✓
Attention Distillation (Pelosin et al., 2022)	ViT	✓	✓
Prompt (Wang et al., 2022b)	ViT	✓	✗

Table 5: Comparison with published state-of-the-art approaches.

aware method for preventing updates to weights that are most useful for previous tasks. Similarly, Masana et al. (2021) propose a ternary-mask method without the need for pretraining or exemplars. Despite their effectiveness, these activation masking methods are task-aware. Our instead is the first parameter isolation, task-agnostic and exemplar-free approach trained entirely from scratch.

Method	Weights/Activations	Class-incremental	Scratch	Exemplar-free
Piggyback (Mallya et al., 2018)	Weights	✗	✗	✓
PackNet (Mallya & Lazebnik, 2018)	Weights	✗	✓	✓
HAT (Serra et al., 2018)	Activations	✗	✓	✓
Ternary (Masana et al., 2021)	Activations	✗	✓	✓
Ours	Activations	✓	✓	✓

Table 6: Comparison with different parameter-isolation approaches from the state-of-the-art.