

# Online Continual Learning with Contrastive Vision Transformer

Zhen Wang, Liu Liu, Yajing Kong, Jiaxian Guo, and Dacheng Tao

The University of Sydney, Australia

{zwan4121,liu.liu1,ykon9947,jguo5934,dacheng.tao}@sydney.edu.au

**Abstract.** Online continual learning (online CL) studies the problem of learning sequential tasks from an online data stream without task boundaries, aiming to adapt to new data while alleviating catastrophic forgetting on the past tasks. This paper proposes a framework Contrastive Vision Transformer (CVT), which designs a focal contrastive learning strategy based on a transformer architecture, to achieve a better stability-plasticity trade-off for online CL. Specifically, we design a new external attention mechanism for online CL that implicitly captures previous tasks' information. Besides, CVT contains learnable focuses for each class, which could accumulate the knowledge of previous classes to alleviate forgetting. Based on the learnable focuses, we design a focal contrastive loss to rebalance contrastive learning between new and past classes and consolidate previously learned representations. Moreover, CVT contains a dual-classifier structure for decoupling learning current classes and balancing all observed classes. The extensive experimental results show that our approach achieves state-of-the-art performance with even fewer parameters on online CL benchmarks and effectively alleviates the catastrophic forgetting.

**Keywords:** Online continual learning, Vision Transformer, Supervised contrastive learning

## 1 Introduction

One of the major challenges in research on artificial neural networks is developing the ability to accumulate knowledge over time from a non-stationary data stream [7,29,38]. Although most successful deep learning techniques can achieve excellent results on pre-collected and fixed datasets, they are incapable of adapting their behavior to the non-stationary environment over time [18,75]. When streaming data comes in continuously, training on the new data can severely interfere with the model's previously learned knowledge of past data, resulting in a drastic drop in performance on the previous tasks. This phenomenon is known as *catastrophic forgetting* or *catastrophic interference* [54,61]. Continual learning (also known as lifelong learning or incremental learning) [62,63,79,75,80,95,18] aims to solve this problem by maintaining and accumulating the acquired knowledge over time from a stream of non-stationary data.




Fig. 1. The overall framework of Contrastive Vision Transformer (CVT).

Continual learning requires neural networks to be both stable to prevent forgetting as well as plastic to learn new concepts, which is referred to as the *stability-plasticity* dilemma [55,27]. The early works have focused on the Task-aware protocol, which selects the corresponding classifier using oracle knowledge of the task identity at inference time [2,46,65,69,18]. For example, regularization-based methods penalize the changes of important parameters when learning new tasks and typically assign a separate classifier for each task [12,42,64,85,93]. Recent studies have focused on a more practical Task-free protocol, which evaluates the network on all classes observed during training without requiring the identity of the task [3,81,8,95,37,90,13]. Among them, rehearsal-based methods that store a small set of seen examples in a limited memory for replaying have demonstrated promising results [4,62,9,76]. This paper focuses on a more realistic and challenging setting: online continual learning (online CL) [34,50,67], where the model learns a sequence of classification tasks with a single pass over the data and without task boundaries. There are also Task-free and Task-aware protocols in online CL.

Inspired by recent breakthroughs in self-supervised learning [16,11,30,72], we find that the knowledge learned by supervised contrastive learning [31,41,28] reveals greater robustness and transferability. The general and transferable knowledge is the essence of what online CL seeks to explore, which could effectively help mitigate forgetting. Unfortunately, it is challenging to employ contrastive learning in the continual setting. There are two main issues as follows: 1) Contrastive learning requires informative negative samples to learn and distinguish different clusters. However, in online CL, previous task data is unavailable or very limited, which causes severe imbalanced contrast between past and new classes. 2) The contrastively learned knowledge may suffer from forgetting since the distribution of the data stream is continually changing.

Considering the superior modeling capability of Vision Transformers [21] recently demonstrated on computer vision tasks [57,84,26,92,17], we take the utilization of the potential of attention mechanism [74] to develop online CL. Overall, we strategically integrate contrastive learning and transformer to model the online data stream. We propose a novel framework, Contrastive Vision Transformer (CVT), to alleviate the forgetting problem and tackle the above imbalance issue of contrastive learning in online CL. An overview of the framework is illustrated in Fig. 1. Specifically, we newly design an effective and efficient trans-

former architecture with external attention to implicitly capture previous tasks' information and reduce the number of parameters. CVT contains learnable focuses for each class, which could accumulate the knowledge of previous classes to alleviate forgetting. Based on the learnable focuses, we design a focal contrastive loss at the attention level to rebalance contrastive learning between new and past classes and improve the inter-class distinction and intra-class aggregation. Moreover, CVT contains a dual-classifier structure: an *injection classifier* is used to inject representation of stream data into the model, mitigating interference with previous knowledge; and an *accumulation classifier* focuses on integrating the previous and new knowledge in a balanced manner.

We systematically compare state-of-the-art and well-established methods for the online CL problem in both the Task-free and Task-aware protocols. Experimental results show that the CVT framework significantly outperforms other approaches in terms of accuracy and forgetting, even with fewer parameters. Ablation study validates each component of the proposed framework.

The main contributions of this paper are three-fold:

- We propose a novel framework Contrastive Vision Transformer (CVT) to achieve a better stability-plasticity trade-off for online CL. CVT contains class-wise learnable focuses, which can accumulate the knowledge of previous classes to alleviate forgetting.
- We design a focal contrastive loss to rebalance contrastive learning between new and past classes and learn more robust representations.
- The extensive experimental results show that CVT achieves state-of-the-art performance with even fewer parameters on online continual learning benchmarks.

## 2 Related Works

### 2.1 Continual Learning Methods

Continual learning (CL) methods have been developed to alleviate catastrophic forgetting in neural networks. These methods can be divided into three main categories: expansion-based, regularization-based, and rehearsal-based methods.

As new tasks arrive, expansion-based methods dynamically expand networks and keep sub-networks related to previous tasks fixed [69,89,1,59,86,45,82,53]. However, most expansion-based methods require task identity during inference in order to allocate distinct sets of parameters to distinct tasks. Regularization-based methods limit the changes in important parameters during the learning of new tasks by estimating the importance of each network parameter for prior tasks [12,42,2,64,43,66,93]. These works differ in how they compute the importance of network parameters.

Rehearsal-based methods [8,37,81,60,15,22,49,79,68,52] alleviate catastrophic forgetting by replaying a subset of data of past tasks stored in limited buffer. iCaRL [62] trains a nearest-class-mean classifier while limiting the change of the representation in later tasks through a self-distillation loss. In addition to

replaying past experiences, HAL [13] keeps predictions on some *anchor* points by an additional objective. IExpressNet [95] introduces representative expression memory while employing a novel Center-expression-distilled loss and having shown satisfactory performance for the task of facial expression recognition. DER++ [9] combines rehearsal with distillation loss [36,83,47] to retrain past experience and obtain state-of-the-art performance. RM [6] proposes an uncertainty-based sampling approach by using uncertainty and data augmentation to improve rehearsal. The proposed method in this paper belongs to the rehearsal-based method.

## 2.2 Online Continual Learning

Online continual learning (online CL) is a more realistic [50,78,87,96,67,49] and difficult setup [34], where the model learns from a non-i.i.d data stream online, without the help of task identifiers or task boundaries both at training and inference stages. Online CL methods are mainly based on rehearsal.

Experience Replay (ER) [63] employs reservoir sampling for memory management and jointly optimizes a network by mixing memory data with online stream data. ERT [10] improves ER by balanced sampling and bias control. AGEM [14] and GEM [49] use episodic memory to compute past task gradients to constrain the online update step. GSS [4] presents a gradient-based sampling to store diversified data for learning more information. ASER [67] is based on the Shapley Value theory to improve the memory buffer update and sampling. CLS [5] uses two extra models to maintain long-term and short-term semantic memories for knowledge consolidation. SCR [52] uses supervised contrastive loss for representation learning and employs the nearest class mean to classify.

## 2.3 Vision Transformers

Transformer model is firstly applied into machine translation tasks [74], and then, Transformers have become the state-of-the-art models for most natural language processing tasks [20,58,24,51,25]. Attention modules are the core components of transformers, which aggregate information from the entire input sequence. Recently, Vision Transformer (ViT) [21] is proposed to makes Transformer architecture scalable for image classification when the data is large enough. Since then, a lot of effort has been dedicated to improving Vision Transformers' data efficiency and model efficiency [92,32,91,40], where an effective direction is to strategically integrate properties of convolution into the Transformer architecture [73,91,48,84,26,94,23,17]. CoAT [88] proposes a conv-attention module that realizes relative position embeddings with convolutions. LeViT [26] builds pyramid attention structure with pooling to learn convolutional-like features. By leveraging sequence pooling and convolutions, CCT [33] eliminates the need for class tokens and positional embeddings.

Nevertheless, current vision transformers may not be applicable to modeling the online data stream, and existing continual learning algorithms developed for CNNs may not be ideal for vision transformers as well. To this end, we propose

a lightweight Contrastive Vision Transformer (CVT) with a focal contrastive loss for online continual learning and achieve better performance than other transformers and CNN baselines.

### 3 Preliminary

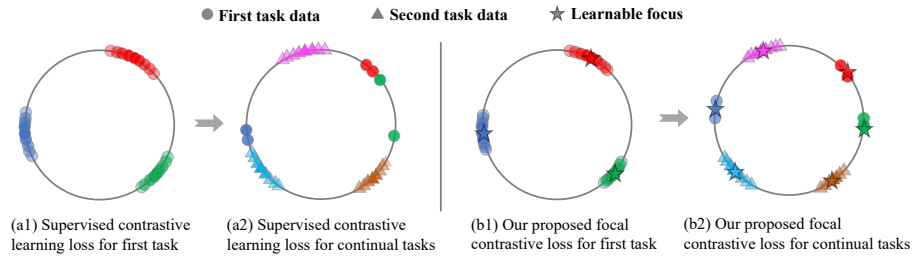
**Problem Setup.** Formally, an online continual learning problem is split in a sequence of  $T$  supervised learning tasks  $\mathcal{T}_t$ ,  $t \in \{1, \dots, T\}$ , where  $T$  is the total number of tasks. Let  $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_T\}$  be the corresponding online data stream, where  $\mathcal{D}_t$  is the dataset of task  $\mathcal{T}_t$ . For task  $\mathcal{T}_t$ , input samples  $x \in \mathcal{X}_t$  and the corresponding ground truth labels  $y \in \mathcal{Y}_t$  are drawn from the i.i.d. distribution  $\mathcal{D}_t$ . A mini-batch of training data  $\mathcal{B}$  from  $\mathcal{D}$  comes gradually in an online stream (each sample is seen only once). Besides, a limited memory buffer  $\mathcal{M}$  saves a small set of training data of seen tasks. The model is trained on  $\mathcal{B} \cup \mathcal{M}$  at each iteration. At task  $\mathcal{T}_t$ , The label space of the model is all observed classes  $\cup_{i=1}^t \mathcal{Y}_i$ , and the model is expected to predict well on all classes at the inference stage.

**Supervised Contrastive Learning (SCL).** SCL [41,28,16] aims to push the representation of samples with different classes farther apart while tightly clustering representation of samples with the same class. Suppose that the classification model can be decomposed into two components: an encoder  $f$  and a classifier  $w$ . Encoder  $f$  maps an image sample  $\mathbf{x}$  to a vectorial embedding (representation)  $\mathbf{z} = f(\mathbf{x})$ . Classifier  $w$  maps the representation  $\mathbf{z}$  to a classification vector  $\hat{\mathbf{y}} = w(\mathbf{z})$ . Without training  $w$ , SCL focuses on training  $f$  as follows: given a batch of  $b$  samples  $\mathcal{B} = \{(\mathbf{x}_k, y_k)\}_{k=1}^b$ , SCL first generates an augmented batch  $\tilde{\mathcal{B}} = \{(\tilde{\mathbf{x}}_k, \tilde{y}_k)\}_{k=1}^{2b}$  by making two random augmentations of  $\mathcal{B}$ , with  $y_k = \tilde{y}_{2k-1} = \tilde{y}_{2k}$ . The SCL loss takes the following form:

$$\mathcal{L}_{\text{SCL}} = \sum_{i \in \mathcal{I}} \frac{-1}{|\mathcal{P}(i)|} \sum_{\mathbf{z}_p \in \mathcal{P}(i)} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau)}{\sum_{\mathbf{z}_j \in \mathcal{A}(i)} \exp(\mathbf{z}_i \cdot \mathbf{z}_j / \tau)}, \quad (1)$$

where  $\mathcal{I}$  represents the set of indices of  $\tilde{\mathcal{B}}$ ;  $\mathcal{A}(i) \equiv \{\mathbf{z}_i : i \in \mathcal{I} \setminus \{i\}\}$  is the set of representations of samples in  $\tilde{\mathcal{B}}$  except for that of  $\mathbf{x}_i$ ;  $\mathcal{P}(i) := \{\mathbf{z}_p \in \mathcal{A}(i) : \tilde{y}_p = \tilde{y}_i\}$  is the set of representations of positive samples (i.e., samples with the same labels) with respect to the anchor  $\tilde{\mathbf{x}}_i$ ;  $\tau \in \mathbb{R}^+$  is a temperature hyperparameter;  $|\mathcal{P}(i)|$  is its cardinality.

Although SCL could learn the transferable representation to help prevent forgetting in online CL, SCL will face new challenges: 1) previous task data is unavailable or very limited due to the streaming fashion, which causes severe imbalanced contrast between past and new classes; 2) the contrastively learned knowledge may suffer from forgetting since the distribution of the data stream is continually changing. For visualization clarity we use a 2D feature space. As illustrated in Fig. 2(a1) and Fig. 2(a2), the imbalanced data stream in online CL makes the representation of previous tasks drift and difficult to be accurately clustered by SCL.



**Fig. 2.** Continually arriving samples are clustered on the hypersphere by contrastive learning (for visualization clarity, we use a 2D representation space). Different colors represent different classes. Supervised contrastive learning loss on online stream data fails to obtain good inter-class distinction and intra-class aggregation caused by the class imbalance. Focal contrastive loss effectively mitigates class imbalance in online CL and accumulates class-wise knowledge by the learnable *focuses*.

## 4 Methodology


To alleviate the forgetting problem in online continue learning, we propose a framework Contrastive Vision Transformer (CVT), which designs a new focal contrastive learning strategy based on the transformer architecture. An overview of the framework is depicted in Fig. 1. CVT plays the strengths of the attention mechanism in online CL, which design an effective transformer architecture with external attention. We tackle the imbalance issue of SCL in online CL by proposing a focal contrastive loss at the attention level. The learnable focuses in CVT can accumulate class-specific knowledge to alleviate the forgetting of previous tasks. Besides, a dual-classifier is used to decouple learning current classes and balancing all seen classes, improving the stability-plasticity trade-off.

In the following, we will go through the description of CVT in terms of both model architecture (Sec. 4.1) and focal contrastive continual learning strategy (Sec. 4.2), respectively.

### 4.1 Model Architecture

Fig. 3 illustrates the CVT architecture. The major contributing components in the architecture include 1) *external attention*, which implicitly captures previous tasks’ information and reduces the number of parameters, and 2) *learnable focuses*, which could maintain and accumulate the knowledge of previous classes.

**External Attention.** CVT plays the strengths of the attention mechanism in online CL. Unlike the vanilla self-attention in vision transformers [21,26] that derives the attention map by computing similarity between self-queries and self-keys [74], we introduce an external attention mechanism [79] to obtain the attention map by computing the affinities between self-queries and a learnable



**Fig. 3.** The architecture of Contrastive Vision Transformer (CVT). CVT architecture is composed of stacked transformer blocks after a simple convolutional block. Shrink module performs downsampling to reduce the resolution of the activation maps and increase their number of channels between CVT stages. Focuses are a set of learnable attention vectors. After a projection layer, two classifiers serve for knowledge injection and accumulation, respectively.

external key  $K_W$  with an attention bias  $B$ , which implicitly injects previous task information in the attention mechanism. Moreover, the proposed architecture can save the number of parameters compared to self-attention.

Let the input tensor be  $X$ , we apply linear transformation with weights  $W_q$  and  $W_v$  to get the vanilla self-query  $Q_X=W_qX$  and self-value  $V_X=W_vX$ , respectively. We employ a linear layer  $K_W$  to replace the input-depended self-key, and explicitly add a learnable attention bias  $B$  to attention maps. Consider  $H$  attention heads, which are uniformly split into  $H$  segments  $Q_X^h, K_W^h, V_X^h$ , and  $B^h$ . The external attention mechanism computes the head-specific attention map  $A^h$  and concatenates the multi-head attention as follows:

$$A^h = \text{Softmax} \left( \frac{\text{Norm}(Q_X^h (K_W^h)^\top) + B^h}{\sqrt{d/H}} \right), \quad X_{out}^h = A^h V_X^h, \quad h = 1, \dots, H, \quad (2)$$

where  $\text{Norm}()$  denotes batch normalization;  $d$  is the dimension of the key.

**Learnable Focuses.** CVT contains learnable *focuses* for each class, which could maintain and accumulate the knowledge of previous classes to alleviate forgetting for online CL. The class-wise learnable focuses  $\mathcal{O} = \{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_C\}$  is a set of learnable attention vectors, as shown in Fig. 3, where a focus  $\mathbf{o}_c$  corresponds to class  $c$ , and  $C$  is the number of the seen classes. The size of the learnable focuses is negligible in relation to the overall model.

When a new class  $c$  appears in the data stream, the corresponding focus  $\mathbf{o}_c$  to class  $c$  starts to participate in the training (refer to Eq. (3)). Even if class  $c$  no longer appears in the data stream afterward, the focus  $\mathbf{o}_c$  always participates in the online CL training and serves as a negative sample for the other classes, as illustrated in Fig. 2(b1) and Fig. 2(b2). Thus, focuses  $\mathcal{O}$  preserve and accumulate the previously learned class-specific knowledge and acts as a forgetting mitigation in online continual learning.

## 4.2 Focal Contrastive Continual Learning

We propose a rehearsal-based focal contrastive learning scheme to 1) tackle the imbalance issue of SCL in online CL and 2) accumulate class-specific knowledge, to alleviate the interference with previous tasks. The learning scheme includes two losses: a focal contrastive loss and a dual-classifier loss, as following.

**Focal Contrastive Loss.** For learning representation continually, we propose a focal contrastive loss  $\mathcal{L}_{\text{FC}}$  in online CL. As mentioned in Sec. 3, during the training phase, the model observes a mini-batch  $\mathcal{B}$  at a time sampled from task  $\mathcal{T}_t$  in the data stream  $\mathcal{D}$ . An input batch for the model is composed of  $\mathcal{B}$  and  $\mathcal{B}_{\mathcal{M}}$  sampled from the memory buffer  $\mathcal{M}$ . The input batch and its augmented view are encoded by CVT blocks to generate attention  $\mathbf{z}$ , as shown in Fig. 3. As mentioned previously, a set of class-wise learnable *focuses*  $\mathcal{O} = \{\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_C\}$  is utilized by focal contrastive loss  $\mathcal{L}_{\text{FC}}$ , where a focus  $\mathbf{o}_c$  is a learnable attention vector for class  $c$  and  $C$  is the number of the seen classes. The FC loss function is defined as:

$$\mathcal{L}_{\text{FC}} = \sum_{i \in \mathcal{I}} \frac{-1}{|\mathcal{P}(i) \cup \mathbf{o}_{\bar{y}_i}|} \sum_{\mathbf{z}_p \in \mathcal{P}(i) \cup \mathbf{o}_{\bar{y}_i}} \delta_{\mathbf{z}_p} \log \frac{\exp(\mathbf{z}_i \cdot \mathbf{z}_p / \tau)}{\sum_{j \in \mathcal{A}(i) \cup \mathcal{O}} \exp(\mathbf{z}_i \cdot \mathbf{z}_j / \tau)}, \quad (3)$$

where

$$\delta_{\mathbf{z}_p} = \begin{cases} \mu, & \mathbf{z}_p \in \mathcal{O} \\ 1.0, & \mathbf{z}_p \in \mathcal{P}(i) \end{cases}, \quad (4)$$

and  $\mu$  is the weight of focuses. We set  $\mu > 1$  to make focuses play a more important role in contrastive learning;  $\mathcal{P}(i)$  and  $\mathcal{A}(i)$  are the same with supervised contrastive learning in Eq. (1);  $\tau \in \mathbb{R}^+$  is a temperature hyperparameter.

The benefits of using focal contrastive loss are two-fold. First, it alleviates the imbalance issue in online CL by employing the class-wise focuses. Second, it accumulates the previous knowledge by the learnable focuses which will continually serve as the prototypes of classes to maintain class-specific information. With a proposed focal contrastive loss  $\mathcal{L}_{\text{FC}}$  in training, CVT rebalances contrastive learning between new and past classes and improve the inter-class distinction and intra-class aggregation, as illustrated in Fig. 2. In Sec. 5.3, we empirically observe that  $\mathcal{L}_{\text{FC}}$  outperforms the original  $\mathcal{L}_{\text{SCL}}$  and boost the online CL.

**Dual-classifier Loss.** We propose a dual-classifier structure to decouple learning current classes and balancing all seen classes, which contains an *injection classifier* to inject new task representation into the model, alleviating interference to previously learned knowledge, and an *accumulation classifier* to integrate past and new knowledge in a balanced manner.



Let  $g(\mathbf{x})$  be the representation of a sample  $\mathbf{x}$  outputted from the Projection of CVT before the classifier. When new data stream batch arrives, we utilize the output from an independent injection classifier to compute a classification loss:

$$\mathcal{L}_I = \sum_{(\mathbf{x}, y) \in \mathcal{B}} \ell(y, f_I(g(\mathbf{x}))), \quad (5)$$

where  $f_I$  denotes the injection classifier and  $\ell$  adopts a cross-entropy loss.  $f_I$  is only trained on stream data and does not participate in the inference stage.

Besides, we employ an accumulation classifier to focus on improving the stability-plasticity trade-off by integrating previous and new knowledge in a balanced manner. The accumulation classifier is used at the inference stage for outputting the prediction. Rehearsing the limited memory data during learning new tasks is a crucial way to maintain previous knowledge. We can replay the exemplars stored in the memory buffer with their ground truth labels. In addition, the accumulation classifier also needs a supervised signal from current task data. Therefore, we give the accumulation classifier loss:

$$\mathcal{L}_A = \alpha \mathbb{E}_{(\mathbf{x}', y') \sim \mathcal{M}} [\ell(y', f_A(g(\mathbf{x}')))] + \beta \sum_{(\mathbf{x}, y) \in \mathcal{B}} \ell(y, f_A(g(\mathbf{x}))), \quad (6)$$

where  $f_A$  denotes the accumulation classifier;  $\alpha$  and  $\beta$  are the coefficients balancing knowledge consolidation. We approximate the expectation by computing gradients on batches sampled from the memory buffer.

Overall, the total loss used in CVT is the sum of Eq. (3), Eq. (5), and Eq. (6):

$$\mathcal{L} = \mathcal{L}_A + \mathcal{L}_I + \gamma \mathcal{L}_{FC}, \quad (7)$$

where  $\gamma$  is the coefficient balancing  $\mathcal{L}_{FC}$ . After updating the whole CVT model on a mini-batch  $\mathcal{B}$ , the memory buffer  $\mathcal{M}$  will be updated with  $\mathcal{B}$  by Reservoir sampling [77,9].

## 5 Experiment

### 5.1 Experimental Setup and Implementation

We consider a strict evaluation setting [38,75] for online continual learning, including task-aware protocol [56] and task-free protocol [39]. For task-aware protocol [56], the task identities are required for each time evaluation. For task-free protocol [39], the task identities are unavailable at inference time.

**Datasets.** Online continual learning benchmarks evaluate the capacity of an algorithm to learn on not independent and identically distributed (non-iid) data. CIFAR-100 [44] contains 100 classes and each class has 100 testing and 500 training color images. TinyImageNet [70] consists 200 classes that include 100,000 images for training and 10,000 images for testing. ImageNet100 [62] contains 100 classes randomly chosen from ILSVRC [19], including about 120,000 images for training and 5,000 images for validation.

Memory Buffer	Method	#Paras	10 splits		20 splits	
			<i>Task-free</i>	<i>Task-aware</i>	<i>Task-free</i>	<i>Task-aware</i>
–	SGD	11.2	5.77±0.35	37.44±1.83	3.53±0.12	36.81±2.63
500	ER [63]	11.2	15.59±1.20	59.72±0.83	12.51±0.77	62.72±1.70
	GEM [49]	11.2	14.34±1.63	50.39±0.28	5.98±0.33	57.15±1.77
	AGEM [14]	11.2	6.35±0.13	39.18±0.27	3.62±0.08	39.55±0.17
	iCaRL [62]	11.2	15.18±0.31	48.95±0.33	12.79±0.26	60.53±0.38
	FDR [8]	11.2	5.97±0.20	32.18±1.60	3.60±0.07	39.98±1.32
	GSS [4]	11.2	10.91±0.36	59.10±0.26	6.33±0.30	62.80±2.31
	DER++ [9]	11.2	15.72±1.29	54.45±1.59	11.29±0.25	63.62±1.05
	HAL [13]	22.4	10.51±0.63	33.70±1.37	7.09±0.39	52.89±2.36
	ERT [10]	11.2	16.28±0.73	60.11±2.18	17.92±0.42	68.08±0.37
	ASER <sub>μ</sub> [67]	11.2	12.42±0.68	53.77±1.30	9.63±0.51	58.91±1.62
	RM [6]	11.2	14.32±0.49	58.76±1.82	13.73±0.60	64.73±2.01
	CLS [5]	33.8	15.06±0.57	59.82±1.24	14.84±0.83	65.74±1.94
	<b>CVT (ours)</b>	<b>8.9</b>	<b>24.45±0.62</b>	<b>62.52±1.43</b>	<b>21.81±0.52</b>	<b>71.23±1.40</b>
1000	ER [63]	11.2	20.41±1.46	63.39±1.37	17.02±1.63	68.52±1.09
	GEM [49]	11.2	16.49±1.08	52.30±0.51	8.40±1.05	62.59±1.89
	AGEM [14]	11.2	6.57±0.12	40.38±0.30	3.74±0.05	42.39±0.42
	iCaRL [62]	11.2	16.31±0.26	50.49±0.18	13.03±0.26	61.13±0.20
	FDR [8]	11.2	6.58±0.31	36.99±0.45	3.72±0.04	42.45±0.39
	GSS [4]	11.2	12.38±0.59	60.75±0.28	7.40±0.23	66.06±1.25
	DER++ [9]	11.2	21.27±1.69	61.80±1.24	13.42±0.50	71.26±0.61
	HAL [13]	22.4	11.81±0.79	39.67±2.63	13.14±0.72	60.03±1.20
	ERT [10]	11.2	23.43±0.58	62.25±1.33	24.58±0.36	72.61±0.67
	ASER <sub>μ</sub> [67]	11.2	14.38±0.43	58.91±1.76	12.79±0.60	62.47±1.82
	RM [6]	11.2	22.41±1.28	61.82±2.13	18.91±1.15	67.30±1.34
	CLS [5]	33.8	19.73±1.17	62.54±1.52	17.06±1.47	70.08±0.85
	<b>CVT (ours)</b>	<b>8.9</b>	<b>28.83±0.86</b>	<b>65.86±1.24</b>	<b>28.15±0.51</b>	<b>75.76±0.93</b>

**Table 1.** Results (overall accuracy %) on CIFAR100 benchmark which is averaged over five runs. #Paras means the number of parameters in the model, which is counted by million.

**Baselines.** We compare CVT with state-of-the-art and well-established Online CL baselines, including 11 rehearsal-based methods (ER [63], GEM [49], AGEM [14], GSS [4], FDR [8], HAL [13], ASER<sub>μ</sub> [67], ERT [10], RM [6], SCR [52], and CLS [5]), 2 methods leveraging Knowledge Distillation (iCaRL [62] and DER++ [9]). Besides, we also compare vision transformers (ViT [21], LeViT [26], CoAT [88], and CCT [33]) with rehearsal strategy for continual learning. We also provide the results of simply performing *SGD* without any countermeasure to alleviate forgetting.

**Metrics.** We evaluate online CL methods in terms of accuracy and forgetting following [12,14,9]. The accuracy is defined by  $\mathbf{A}_T = \frac{1}{T} \sum_{t=1}^T a_{T,t}$ , and the forgetting is defined by  $\mathbf{F}_T = \frac{1}{T-1} \sum_{t=1}^{T-1} \max_{i \in \{1, \dots, T-1\}} (a_{i,t} - a_{T,t})$ , where  $a_{T,t}$  is the inference accuracy on task  $\mathcal{T}_t$  when the model finished learning task  $\mathcal{T}_T$ .

Memory Buffer	Method	#Paras	TinyImageNet		#Paras	ImageNet100	
			Task-free	Task-aware		Task-free	Task-aware
–	SGD	11.2	4.54±0.03	26.25±0.16	11.2	3.98±0.02	19.05±0.17
500	ER [63]	11.2	9.71±0.18	42.76±0.35	11.2	9.88±0.46	32.38±0.63
	AGEM [14]	11.2	4.63±0.08	27.86±0.13	11.2	3.38±0.08	21.80±0.15
	iCaRL [62]	11.2	6.17±1.03	27.22±1.52	11.2	7.70±0.32	20.45±0.80
	FDR [8]	11.2	5.19±0.18	28.23±0.46	11.2	3.34±0.16	19.24±0.07
	DER++ [9]	11.2	9.56±0.69	40.52±0.47	11.2	10.30±0.24	29.20±0.38
	ERT [10]	11.2	9.95±0.72	40.42±1.57	11.2	10.28±0.25	28.53±0.48
	ASER <sub>μ</sub> [67]	11.2	9.22±0.25	41.09±0.68	11.2	9.75±0.57	31.71±0.79
	RM [6]	11.2	8.39±1.37	41.63±0.74	11.2	8.53±0.68	28.30±0.52
	SCR [52]	11.2	9.08±0.74	39.85±0.93	11.2	8.81±0.79	29.62±1.24
	<b>CVT (ours)</b>	<b>9.0</b>	<b>14.71±1.04</b>	<b>43.93±1.42</b>	<b>9.4</b>	<b>14.82±0.31</b>	<b>36.74±0.46</b>
1000	ER [63]	11.2	12.46±0.45	45.50±0.61	11.2	10.42±0.51	34.26±0.43
	AGEM [14]	11.2	4.92±0.13	28.38±0.15	11.2	3.66±0.05	23.56±0.19
	iCaRL [62]	11.2	6.91±0.52	28.56±0.37	11.2	8.93±0.48	22.37±1.04
	FDR [8]	11.2	5.27±0.12	28.94±0.36	11.2	3.58±0.09	21.28±0.11
	DER++ [9]	11.2	12.97±0.42	47.21±0.33	11.2	13.94±0.71	40.02±0.39
	ERT [10]	11.2	13.84±0.77	44.65±0.79	11.2	12.26±0.23	33.88±0.60
	ASER <sub>μ</sub> [67]	11.2	12.26±0.38	46.02±0.82	11.2	11.38±0.54	35.76±1.28
	RM [6]	11.2	11.73±0.89	45.89±0.64	11.2	11.85±1.17	32.72±0.85
	SCR [52]	11.2	10.19±1.14	43.58±1.20	11.2	10.74±0.85	31.84±2.27
	<b>CVT (ours)</b>	<b>9.0</b>	<b>16.54±1.22</b>	<b>48.50±0.88</b>	<b>9.4</b>	<b>18.02±0.25</b>	<b>42.61±0.72</b>

**Table 2.** Results (overall accuracy %) on TinyImageNet and ImageNet100, which are averaged over three runs. #Paras means the number of parameters in the model, which is counted by million.

**Implementation Details.** In order to compare each method fairly, we train all networks using stochastic gradient descent (SGD) optimizer. The images used for training are randomly cropped and flipped for each method following [9,10,68]. We adopt 1 epoch with mini-batch size of 10 for all datasets, following [62,9,10,93]. Online continual learning baselines use ResNet18 [35] as backbone and cross-entropy as the classification loss, following [14,68,6,9,13,71]. The implementation of transformer block is based on LeViT [26] and ViT [21]. CVT framework employs GELU activation and dropout in transformer blocks and uses a global average pooling to the last activation map.

## 5.2 Comparison to State-of-the-Art Methods

**Evaluation on CIFAR100.** Following the setting proposed in [62,89], we train all 100 classes in several splits, including 10, 20 incremental tasks. Table 1 summarizes the overall accuracy on CIFAR100 with 500 and 1000 memory sizes. It is demonstrated that CVT outperforms other baselines by a considerable margin in different incremental splits, e.g., CVT can improve the accuracy of continual learning by more than 8% in 10-split with 500 memory capacity. Especially in the case of small memory, the advantage of CVT is more obvious, which indicates CVT can effectively alleviate the imbalance issue in online CL. It is worth noting that although CVT uses fewer parameters (8.9M) than other methods




Fig. 4. Forgetting results (%) on CIFAR100 (lower is better).

(11.2M~33.8M), it can still achieve superior performance. One reason is that CVT inherits the merits of transformers for modeling the stream of tasks without stacking a lot of parameters, and besides, the number of parameters for the proposed learnable focuses is extremely small.

**Evaluation on ImageNet datasets.** Table 2 summarizes the evaluation results for the TinyImageNet and ImageNet100 datasets with 10 splits. It is demonstrated that CVT consistently surpasses other methods with a considerable margin for Task-free and Task-aware on TinyImageNet and ImageNet100 datasets. Specifically, our method outperforms the state-of-the-art with about **4.3%** for the Task-aware accuracy on the ImageNet100 benchmark. For TinyImageNet benchmark, the Task-free accuracy is improved from 9.95% to 14.71%(+**4.76%**). Moreover, CVT takes fewer parameters compared to other CNN-based methods.

**Forgetting.** To compare the alleviating forgetting capability, we assess the *average forgetting* [9,12] that measures the performance degradation in subsequent tasks. As shown in Fig. 4, CVT suffers from less forgetting than all the other baselines in both of Task-free and Task-aware settings with memory buffer 1000 on CIFAR100. This is because CVT utilizes the focal contrastive loss and dual-classifier loss, which improve the stability of the vision transformer network.

**Incremental Performance.** We also evaluate the *average incremental performance* [62,9] under the Task-free protocol with 500 memory buffer, which is the result of evaluating on all the tasks observed so far after completing each task. As illustrated in Fig. 5, the results are curves of accuracy and forgetting after each task. It is observed that during the learning process, most methods degrade rapidly as new tasks arrive, while our method consistently outperforms the state-of-the-art methods in both accuracy and forgetting.

### 5.3 Ablation Study and Analysis

**Comparison to Transformer and CNN Backbone.** We compare CVT to Vision Transformer networks (ViT [21], LeViT [26], CoAT [88], and CCT [33])

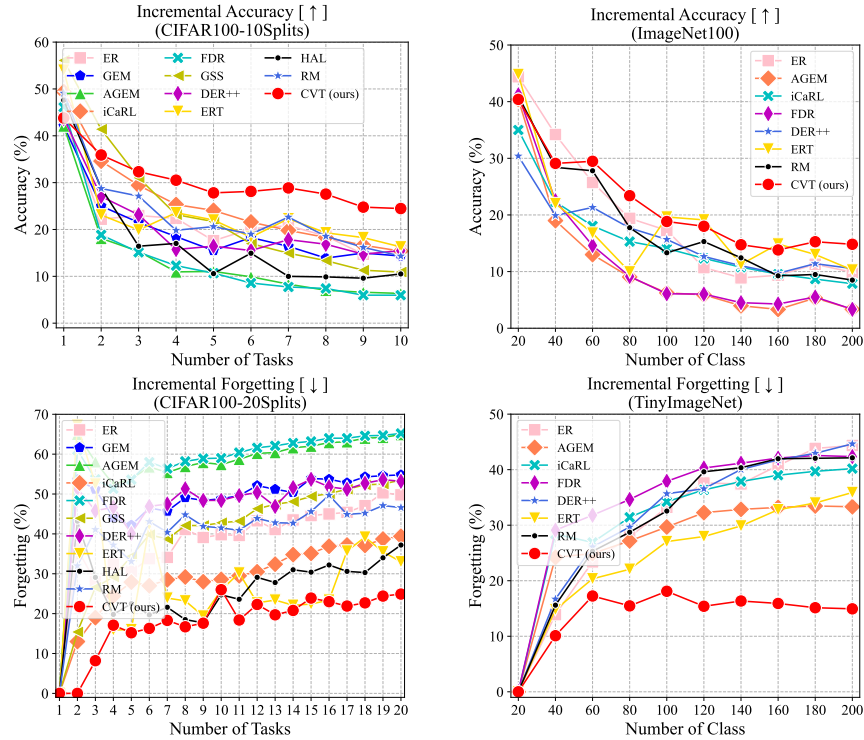


Fig. 5. Incremental performance evaluated on tasks observed so far. [↑] higher is better, [↓] lower is better (best seen in color).

and the CNN benchmark ResNet18 [35] under the proposed rehearsal strategy in online continual learning. Table 3 demonstrates the results of accuracy and forgetting on CIFAR100 and TinyImageNet with 500 memory. It is observed that ViT is not up to the task of online continual learning, since it is “data-hungry” and only fits i.i.d. large datasets. Besides, LeViT, CoAT, and CCT contain CNN structures to obtain inductive biases, which still suffer from catastrophic forgetting in online CL. We can find that using Vision Transformer directly for online CL can not consistently outperform CNN-based networks. Our proposed CVT architecture essentially inherits the merits of CNN and transformers and thus works well in online streaming data and modeling long-dependencies in the input data. Moreover, CVT even takes fewer parameters to achieve better performance for online CL, which also benefits from the focal contrastive loss and the dual-classifier structure.

**Effect of Each Component.** To assess the effects of the components in CVT, we perform ablation study in terms of accuracy and forgetting. From Table 3 we can observe that the proposed focal contrastive loss  $\mathcal{L}_{FC}$  plays an important role in alleviating catastrophic forgetting and accumulating knowledge. However,

Method	#Paras	CIFAR100		#Paras	TinyImageNet	
		Accuracy[↑]	Forgetting[↓]		Accuracy[↑]	Forgetting[↓]
ViT [21]	16.2	8.48	35.56	16.3	7.91	42.26
LeViT [26]	10.9	14.55	44.53	12.1	9.02	41.41
CoAT [88]	10.3	13.17	47.64	11.3	8.78	39.80
CCT [33]	<b>4.3</b>	13.86	51.06	<b>4.4</b>	9.97	40.21
ResNet18 [35]	11.2	15.72	43.82	11.2	9.56	42.13
ResNet18 + $\mathcal{L}_{FC}$	11.3	17.49	38.73	11.3	10.17	36.49
ResNet18 + dual-classifier	11.2	18.84	35.38	11.2	10.91	33.55
CVT - $\mathcal{L}_{FC}$	8.8	19.92	26.16	8.9	12.47	19.43
CVT - $\mathcal{L}_{FC}$ + $\mathcal{L}_{SCL}$	8.9	20.73	25.52	9.0	12.59	20.47
CVT - dual-classifier	8.9	22.08	29.81	9.0	13.63	18.95
CVT (ours)	8.9	<b>24.45</b>	<b>21.86</b>	9.0	<b>14.71</b>	<b>16.32</b>

**Table 3.** Ablation study on backbone and each component of CVT. “-” indicates the removal operation. “+” represents an add component operation.

if we simply use supervised contrastive learning loss  $\mathcal{L}_{SCL}$  to replace  $\mathcal{L}_{FC}$ , we find that the forgetting problem is not mitigated compared to not using any contrastive loss. This is because using  $\mathcal{L}_{SCL}$  directly could cause a severe imbalance between new and past classes in online CL, which limits the learning of transferable representation. While  $\mathcal{L}_{FC}$  can overcome the issue by utilizing the learnable focuses to boost the performance of online CL. This supports that  $\mathcal{L}_{FC}$  can **rebalance contrastive learning between new and past classes and improve the inter-class distinction and intra-class aggregation**. Besides, we can see the dual-classifier loss obtains 2.37% and 7.97% gain in terms of accuracy and forgetting on CIFAR100, respectively. The results of Table 3 demonstrate the effectiveness of each component of CVT.

## 6 Conclusion

In this paper, we propose a novel attention-based framework, Contrastive Vision Transformer (CVT), to effectively mitigate the catastrophic forgetting for online CL. To the best of our knowledge, this paper is the first in the literature to design a Transformer for online CL. CVT contains external attention and learnable focuses to accumulate previous knowledge and maintain class-specific information. With a proposed focal contrastive loss in training, CVT rebalances contrastive continual learning between new and past classes and improves the inter-class distinction and intra-class aggregation. Moreover, CVT designs a dual-classifier structure to decouple learning current classes and balancing all seen classes. Extensive experimental results show that our approach significantly outperforms current state-of-the-art methods with fewer parameters. Ablation study validates the effectiveness of each proposed component.

**Acknowledgments** This work is supported by ARC FL-170100117, DP-180103424, IC-190100031, and LE-200100049.

## References

1. Abati, D., Tomczak, J., Blankevoort, T., Calderara, S., Cucchiara, R., Bejnordi, B.E.: Conditional channel gated networks for task-aware continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3931–3940 (2020) [3](#)
2. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: Learning what (not) to forget. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 139–154 (2018) [2](#), [3](#)
3. Aljundi, R., Chakravarty, P., Tuytelaars, T.: Expert gate: Lifelong learning with a network of experts. In: CVPR. pp. 3366–3375 (2017) [2](#)
4. Aljundi, R., Lin, M., Goujaud, B., Bengio, Y.: Gradient based sample selection for online continual learning. In: Advances in Neural Information Processing Systems (2019) [2](#), [4](#), [10](#)
5. Arani, E., Sarfraz, F., Zonooz, B.: Learning fast, learning slow: A general continual learning method based on complementary learning system. In: International Conference on Learning Representations (ICLR) (2022) [4](#), [10](#)
6. Bang, J., Kim, H., Yoo, Y., Ha, J.W., Choi, J.: Rainbow memory: Continual learning with a memory of diverse samples. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8218–8227 (2021) [4](#), [10](#), [11](#)
7. Bengio, Y., Goodfellow, I., Courville, A.: Deep learning, vol. 1. MIT press Massachusetts, USA: (2017) [1](#)
8. Benjamin, A.S., Rolnick, D., Kording, K.P.: Measuring and regularizing networks in function space. International Conference on Learning Representations (2019) [2](#), [3](#), [10](#), [11](#)
9. Buzzega, P., Boschini, M., Porrello, A., Abati, D., Calderara, S.: Dark Experience for General Continual Learning: a Strong, Simple Baseline. In: Advances in Neural Information Processing Systems (NeurIPS) (2020) [2](#), [4](#), [9](#), [10](#), [11](#), [12](#)
10. Buzzega, P., Boschini, M., Porrello, A., Calderara, S.: Rethinking experience replay: a bag of tricks for continual learning. In: International Conference on Pattern Recognition (ICPR). pp. 2180–2187. IEEE (2021) [4](#), [10](#), [11](#)
11. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. Advances in Neural Information Processing Systems **33**, 9912–9924 (2020) [2](#)
12. Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 532–547 (2018) [2](#), [3](#), [10](#), [12](#)
13. Chaudhry, A., Gordo, A., Dokania, P.K., Torr, P.H., Lopez-Paz, D.: Using hindsight to anchor past knowledge in continual learning. In: AAAI (2021) [2](#), [4](#), [10](#), [11](#)
14. Chaudhry, A., Ranzato, M., Rohrbach, M., Elhoseiny, M.: Efficient Lifelong Learning with A-GEM. In: International Conference on Learning Representations (ICLR) (2019) [4](#), [10](#), [11](#)
15. Chen, S., Wang, L., Wang, Z., Yan, Y., Wang, D.H., Zhu, S.: Learning meta-adversarial features via multi-stage adaptation network for robust visual object tracking. Neurocomputing **491**, 365–381 (2022) [3](#)
16. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020) [2](#), [5](#)

17. Dai, Z., Liu, H., Le, Q.V., Tan, M.: Coatnet: Marrying convolution and attention for all data sizes. In: *Advances in Neural Information Processing Systems* (2021) [2](#), [4](#)
18. Delange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., Tuytelaars, T.: A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2021) [1](#), [2](#)
19. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 248–255 (2009) [9](#)
20. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018) [4](#)
21. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations* (2021) [2](#), [4](#), [6](#), [10](#), [11](#), [12](#), [14](#)
22. Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E.: Podnet: Pooled outputs distillation for small-tasks incremental learning. In: *European Conference on Computer Vision (ECCV)*. pp. 86–102 (2020) [3](#)
23. Douillard, A., Ramé, A., Couairon, G., Cord, M.: Dytox: Transformers for continual learning with dynamic token expansion. *Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) [4](#)
24. Duan, Y., Wang, Z., Wang, J., Wang, Y.K., Lin, C.T.: Position-aware image captioning with spatial relation. *Neurocomputing* **497**, 28–38 (2022) [4](#)
25. Girdhar, R., Carreira, J., Doersch, C., Zisserman, A.: Video action transformer network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 244–253 (2019) [4](#)
26. Graham, B., El-Nouby, A., Touvron, H., Stock, P., Joulin, A., Jégou, H., Douze, M.: Levit: A vision transformer in convnet’s clothing for faster inference. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 12259–12269 (October 2021) [2](#), [4](#), [6](#), [10](#), [11](#), [12](#), [14](#)
27. Grossberg, S.: Adaptive resonance theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural networks* **37**, 1–47 (2013) [2](#)
28. Gunel, B., Du, J., Conneau, A., Stoyanov, V.: Supervised contrastive learning for pre-trained language model fine-tuning. *arXiv preprint arXiv:2011.01403* (2020) [2](#), [5](#)
29. Guo, J., Gong, M., Liu, T., Zhang, K., Tao, D.: LTF: A label transformation framework for correcting label shift. In: *ICML*. vol. 119, pp. 3843–3853 (2020) [1](#)
30. Guo, J., Li, J., Fu, H., Gong, M., Zhang, K., Tao, D.: Alleviating semantics distortion in unsupervised low-level image-to-image translation via structure consistency constraint. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 18249–18259 (2022) [2](#)
31. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*. vol. 2, pp. 1735–1742. IEEE (2006) [2](#)
32. Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., Tang, Y., Xiao, A., Xu, C., Xu, Y., Yang, Z., Zhang, Y., Tao, D.: A survey on visual transformer. *ArXiv abs/2012.12556* (2020) [4](#)



33. Hassani, A., Walton, S., Shah, N., Abuduweili, A., Li, J., Shi, H.: Escaping the big data paradigm with compact transformers. arXiv preprint arXiv:2104.05704 (2021) [4](#), [10](#), [12](#), [14](#)
34. He, J., Mao, R., Shao, Z., Zhu, F.: Incremental learning in online scenario. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13926–13935 (2020) [2](#), [4](#)
35. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016) [11](#), [13](#), [14](#)
36. Hinton, G., Vinyals, O., Dean, J.: Distilling the Knowledge in a Neural Network. In: NeurIPS workshop (2014) [4](#)
37. Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: CVPR. pp. 831–839 (2019) [2](#), [3](#)
38. Hsu, Y.C., Liu, Y.C., Ramasamy, A., Kira, Z.: Re-evaluating Continual Learning Scenarios: A Categorization and Case for Strong Baselines. In: NeurIPS Continual learning Workshop (2018) [1](#), [9](#)
39. Jin, X., Sadhu, A., Du, J., Ren, X.: Gradient-based editing of memory examples for online task-free continual learning. *Advances in Neural Information Processing Systems* **34** (2021) [9](#)
40. Khan, S., Naseer, M., Hayat, M., Zamir, S.W., Khan, F.S., Shah, M.: Transformers in vision: A survey. arXiv preprint arXiv:2101.01169 (2021) [4](#)
41. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *Advances in Neural Information Processing Systems* **33** (2020) [2](#), [5](#)
42. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**(13), 3521–3526 (2017) [2](#), [3](#)
43. Kong, Y., Liu, L., Wang, J., Tao, D.: Adaptive curriculum learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5067–5076 (2021) [3](#)
44. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009) [9](#)
45. Li, X., Zhou, Y., Wu, T., Socher, R., Xiong, C.: Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In: International Conference on Machine Learning (ICML) (2019) [3](#)
46. Li, Z., Hoiem, D.: Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **40**(12) (2017) [2](#)
47. Liu, X., Wang, X., Matwin, S.: Improving the interpretability of deep neural networks with knowledge distillation. In: 2018 IEEE International Conference on Data Mining Workshops (ICDMW). IEEE (2018) [4](#)
48. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012–10022 (2021) [4](#)
49. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2017) [3](#), [4](#), [10](#)
50. Losing, V., Hammer, B., Wersing, H.: Incremental on-line learning: A review and comparison of state of the art algorithms. *Neurocomputing* **275**, 1261–1274 (2018) [2](#), [4](#)

51. Luong, T., Pham, H., Manning, C.D.: Effective approaches to attention-based neural machine translation. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. pp. 1412–1421. Association for Computational Linguistics (Sep 2015) [4](#)
52. Mai, Z., Li, R., Kim, H., Sanner, S.: Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 3584–3594. IEEE (2021) [3](#), [4](#), [10](#), [11](#)
53. Mallya, A., Davis, D., Lazebnik, S.: Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018) [3](#)
54. McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of Learning and Motivation*, vol. 24, pp. 109–165. Academic Press (1989) [1](#)
55. Mermillod, M., Bugajska, A., Bonin, P.: The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects. *Frontiers in psychology* **4**, 504 (2013) [2](#)
56. Pham, Q., Liu, C., Hoi, S.: Dualnet: Continual learning, fast and slow. *Advances in Neural Information Processing Systems* **34** (2021) [9](#)
57. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. arXiv preprint arXiv:2103.00020 (2021) [2](#)
58. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019) [4](#)
59. Rajasegaran, J., Hayat, M., Khan, S.H., Khan, F.S., Shao, L.: Random path selection for continual learning. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2019) [3](#)
60. Rannen, A., Aljundi, R., Blaschko, M.B., Tuytelaars, T.: Encoder based lifelong learning. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1320–1328 (2017) [3](#)
61. Ratcliff, R.: Connectionist models of recognition memory: constraints imposed by learning and forgetting functions. *Psychological review* **97**(2), 285 (1990) [1](#)
62. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR) (2017) [1](#), [2](#), [3](#), [9](#), [10](#), [11](#), [12](#)
63. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., Tesauro, G.: Learning to Learn without Forgetting by Maximizing Transfer and Minimizing Interference. In: *International Conference on Learning Representations (ICLR)* (2019) [1](#), [4](#), [10](#), [11](#)
64. Schwarz, J., Czarnecki, W., Luketina, J., Grabska-Barwinska, A., Teh, Y.W., Pascanu, R., Hadsell, R.: Progress & Compress: A scalable framework for continual learning. In: *International Conference on Machine Learning* (2018) [2](#), [3](#)
65. Serra, J., Suris, D., Miron, M., Karatzoglou, A.: Overcoming catastrophic forgetting with hard attention to the task. In: *International Conference on Machine Learning*. vol. 80, pp. 4548–4557 (2018) [2](#)
66. Shi, Y., Yuan, L., Chen, Y., Feng, J.: Continual learning via bit-level information preserving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16674–16683 (2021) [3](#)
67. Shim, D., Mai, Z., Jeong, J., Sanner, S., Kim, H., Jang, J.: Online class-incremental continual learning with adversarial shapley value. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 9630–9638 (2021) [2](#), [4](#), [10](#), [11](#)

68. Simon, C., Koniusz, P., Harandi, M.: On learning the geodesic path for incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1591–1600 (2021) [3](#), [11](#)
69. Singh, P., Mazumder, P., Rai, P., Namboodiri, V.P.: Rectification-based knowledge retention for continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15282–15291 (2021) [2](#), [3](#)
70. Stanford: Tiny ImageNet Challenge (CS231n) (2015), <http://tiny-imagenet.herokuapp.com/> [9](#)
71. Tang, S., Chen, D., Zhu, J., Yu, S., Ouyang, W.: Layerwise optimization by gradient decomposition for continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9634–9643 (2021) [11](#)
72. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. arXiv preprint arXiv:1906.05849 (2019) [2](#)
73. Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12894–12904 (2021) [4](#)
74. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017) [2](#), [4](#), [6](#)
75. van de Ven, G.M., Tolias, A.S.: Three continual learning scenarios. NeurIPS Continual Learning Workshop (2018) [1](#), [9](#)
76. Verwimp, E., De Lange, M., Tuytelaars, T.: Rehearsal revealed: The limits and merits of revisiting samples in continual learning. In: ICCV (2021) [2](#)
77. Vitter, J.S.: Random sampling with a reservoir. ACM Transactions on Mathematical Software (TOMS) **11**(1), 37–57 (1985) [9](#)
78. Wang, Z., Duan, Y., Liu, L., Tao, D.: Multi-label few-shot learning with semantic inference. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 15917–15918 (2021) [4](#)
79. Wang, Z., Liu, L., Duan, Y., Kong, Y., Tao, D.: Continual learning with lifelong vision transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 171–181 (June 2022) [1](#), [3](#), [6](#)
80. Wang, Z., Liu, L., Duan, Y., Tao, D.: Continual learning with embeddings: Algorithm and analysis. In: ICML 2021 Workshop on Theory and Foundation of Continual Learning (2021) [1](#)
81. Wang, Z., Liu, L., Duan, Y., Tao, D.: Continual learning through retrieval and imagination. Proceedings of the AAAI Conference on Artificial Intelligence **36**(8), 8594–8602 (2022) [2](#), [3](#)
82. Wang, Z., Liu, L., Duan, Y., Tao, D.: Sin: Semantic inference network for few-shot streaming label learning. IEEE Transactions on Neural Networks and Learning Systems pp. 1–14 (2022) [3](#)
83. Wang, Z., Liu, L., Tao, D.: Deep streaming label learning. In: International Conference on Machine Learning (ICML). vol. 119, pp. 9963–9972 (2020) [4](#)
84. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 22–31 (October 2021) [2](#), [4](#)
85. Xi, H., Aussel, D., Liu, W., Waller, S.T., Rey, D.: Single-leader multi-follower games for the regulation of two-sided mobility-as-a-service markets. European Journal of Operational Research (2022) [2](#)

86. Xi, H., He, L., Zhang, Y., Wang, Z.: Bounding the efficiency gain of differentiable road pricing for evs and gvs to manage congestion and emissions. *PloS one* **15**(7), e0234204 (2020) [3](#)
87. Xi, H., Liu, W., Rey, D., Waller, S.T., Kilby, P.: Incentive-compatible mechanisms for online resource allocation in mobility-as-a-service systems. *arXiv preprint arXiv:2009.06806* (2020) [4](#)
88. Xu, W., Xu, Y., Chang, T., Tu, Z.: Co-scale conv-attentional image transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 9981–9990 (October 2021) [4](#), [10](#), [12](#), [14](#)
89. Yan, S., Xie, J., He, X.: Der: Dynamically expandable representation for class incremental learning. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) [3](#), [11](#)
90. Yu, L., Twardowski, B., Liu, X., Herranz, L., Wang, K., Cheng, Y., Jui, S., Weijer, J.v.d.: Semantic drift compensation for class-incremental learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6982–6991 (2020) [2](#)
91. Yu, P., Chen, Y., Jin, Y., Liu, Z.: Improving vision transformers for incremental learning. *arXiv preprint arXiv:2112.06103* (2021) [4](#)
92. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 558–567 (October 2021) [2](#), [4](#)
93. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: *International Conference on Machine Learning (ICML)* (2017) [2](#), [3](#), [11](#)
94. Zhao, X., Zhu, J., Luo, B., Gao, Y.: Survey on facial expression recognition: History, applications, and challenges. *IEEE MultiMedia* **28**(4), 38–44 (2021) [4](#)
95. Zhu, J., Luo, B., Zhao, S., Ying, S., Zhao, X., Gao, Y.: Iexpressnet: Facial expression recognition with incremental classes. In: *Proceedings of the 28th ACM International Conference on Multimedia*. pp. 2899–2908 (2020) [1](#), [2](#), [4](#)
96. Zhu, J., Wei, Y., Feng, Y., Zhao, X., Gao, Y.: Physiological signals-based emotion recognition via high-order correlation learning. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **15**(3s), 1–18 (2019) [4](#)