

Continual Learning Based on OOD Detection and Task Masking

Gyuhak Kim

University of Illinois at Chicago

gkim87@uic.edu

Changnan Xiao

ByteDance

xiaochangnan@bytedance.com

Sepideh Esmaeilpour

University of Illinois at Chicago

sesmae2@uic.edu

Bing Liu

University of Illinois at Chicago

liub@uic.edu

Abstract

Existing continual learning techniques focus on either task incremental learning (TIL) or class incremental learning (CIL) problem, but not both. CIL and TIL differ mainly in that the task-id is provided for each test sample during testing for TIL, but not provided for CIL. Continual learning methods intended for one problem have limitations on the other problem. This paper proposes a novel unified approach based on out-of-distribution (OOD) detection and task masking, called CLOM, to solve both problems. The key novelty is that each task is trained as an OOD detection model rather than a traditional supervised learning model, and a task mask is trained to protect each task to prevent forgetting. Our evaluation shows that CLOM outperforms existing state-of-the-art baselines by large margins. The average TIL/CIL accuracy of CLOM over six experiments is 87.6/67.9% while that of the best baselines is only 84.4/55.0%. The code of our system is available at <https://github.com/k-gyuhak/CLOM>.

1. Introduction

Continual learning (CL) learns a sequence of tasks $\langle 1, 2, \dots, k, \dots \rangle$ incrementally [12]. Each task k has its dataset $\mathcal{D}_k = \{(\mathbf{x}_k^i, y_k^i)_{i=1}^{n_k}\}$, where $\mathbf{x}_k^i \in \mathbf{X}$ is a data sample in task k and $y_k^i \in \mathbf{Y}_k$ is the class label of \mathbf{x}_k^i and \mathbf{Y}_k is the set of classes of task k . The key challenge of CL is *catastrophic forgetting* (CF) [46], which refers to the situation where the learning of a new task may significantly change the network weights learned for old tasks, degrading the model accuracy for old tasks. Researchers have mainly worked on two CL problems: *class incremental/continual learning* (CIL) and *task incremental/continual learning* (TIL) [14, 66]. The main difference between CIL and TIL is that in TIL, the task-id k is provided for each test sample \mathbf{x} during testing so that only the model for task k is used to classify \mathbf{x} , while

in CIL, the task-id k for each test sample is not provided.

Existing CL techniques focus on either CIL or TIL [12, 50]. In general, CIL methods are designed to function without given task-id to perform class prediction over all classes in the tasks learned so far, and thus tend to forget the previous task performance due to plasticity for the new task. TIL methods are designed to function with a task-id for class prediction within the task. They are more stable to retain previous within-task knowledge, but incompetent if task-id is unknown (see Sec. 4).

This paper proposes a novel and unified method called CLOM (Continual Learning based on OOD detection and Task Masking) to solve both problems by overcoming the limitations of CIL and TIL methods. CLOM has two key mechanisms: (1) a task mask mechanism for protecting each task model to overcome CF, and (2) a learning method for building a model for each task based on out-of-distribution (OOD) detection. The task mask mechanism is inspired by *hard attention* in the TIL system HAT [61]. The OOD detection based learning method for building each task model is quite different from the classic supervised learning used in existing TIL systems. It is, in fact, the key novelty and enabler for CLOM to work effectively for CIL.

OOD detection is stated as follows [4]: Given a training set of n classes, called the *in-distribution* (IND) training data, we want to build a model that can assign correct classes to IND test data and reject or detect OOD test data that do not belong to any of the n IND training classes. The OOD rejection capability makes a TIL model effective for CIL problem because during testing, if a test sample does not belong to any of the classes of a task, it will be rejected by the model of the task. Thus, only the task that the test sample belongs to will accept it and classify it to one of its classes. No task-id is needed. The OOD detection algorithm in CLOM is inspired by several recent advances in self-supervised learning, data augmentation [21], contrastive learning [11, 48], and their applications to OOD detection [18, 23, 65].

The main contributions of this work are as follows:

- It proposes a novel CL method CLOM, which is essentially a TIL system in training, but solves both TIL and CIL problems in testing. Existing methods mainly solve either TIL or CIL problem, but are weak on the other problem.
- CLOM uses task masks to protect the model for each task to prevent CF and OOD detection to build each task model, which, to our knowledge, has not been done before. More importantly, since the task masks can prevent CF, continual learning performance gets better as OOD models get better too.
- CLOM needn't to save any replay data. If some past data is saved for output calibration, it performs even better.

Experimental results show that CLOM improves state-of-the-art baselines by large margins. The average TIL/CIL accuracy of CLOM over six different experiments is 87.6/67.9% while that of the best baseline is only 84.4/55.0%.

2. Related Work

Many approaches have been proposed to deal with CF in CL. Using regularization [33] and knowledge distillation [38] to minimize the change to previous models are two popular approaches [6, 7, 14, 15, 25, 27, 30, 37, 40, 53, 55, 59, 60, 71, 75]. Memorizing some old examples and using them to adjust the old models in learning a new task is another popular approach (called *replay*) [5, 10, 13, 24, 41, 43, 51, 54, 56, 58, 70, 76]. Several systems learn to generate pseudo training data of old tasks and use them to jointly train the new task, called *pseudo-replay* [17, 20, 25, 28, 30, 49, 57, 60, 62, 69, 69]. CLOM differs from these approaches as it does not replay any old task data to prevent forgetting and can function with/without saving some old data. *Parameter isolation* is yet another popular approach, which makes different subsets (which may overlap) of the model parameters dedicated to different tasks using masks [15, 29, 61] or finding a sub-network for each task by pruning [26, 45, 68]. CLOM uses parameter isolation, but it differs from these approaches as it combines the idea of parameter isolation and OOD detection, which can solve both TIL and CIL problems effectively.

There are also many other approaches, *e.g.*, orthogonal projection [19, 32, 74], a network of experts [3] and generalized CL [47], etc. PASS [77] uses data rotation and regularizes them. Co²L [8] is a replay method that uses contrastive loss on old samples. CLOM also uses rotation and contrastive loss, but its CF handling is based on masks. None of the existing methods uses OOD detection.

CLOM is a TIL method that can also solve CIL problems. Many TIL systems exist [15], *e.g.*, GEM [43], A-GEM [10], UCL [2], ADP [72], CCLL [63], Orthog-Subspace [9], HyperNet [67], PackNet [45], CPG [26], SupSup [68], HAT [61], and CAT [29]. GEM is a replay based method and UCL is a regularization based method. A-GEM [10] improves GEM's efficiency. ADP decomposes parameters into shared and adaptive parts to construct an order robust TIL system. CCLL uses task-adaptive calibration on convolution layers. Orthog-Subspace learns each task in subspaces orthogonal to each other. HyperNet initializes task-specific parameters conditioned on task-id. PackNet, CPG and SupSup find an isolated sub-network for each task and use it at inference. HAT and CAT protect previous tasks by masking the important parameters. Our CLOM also uses this general approach, but its model building for each task is based on OOD detection, which has not been used by existing TIL methods. It also performs well in the CIL setting (see Sec. 4.2).

Related work on out-of-distribution (OOD) detection (also called open set detection) is also extensive. Excellent surveys include [4, 16]. The model building part of CLOM is inspired by the latest method in [65] based on data augmentation [21] and contrastive learning [11].

3. Proposed CLOM Technique

As mentioned earlier, CLOM is a TIL system that solves both TIL and CIL problems. It takes the parameter isolation approach for TIL. For each task k (task-id), a model is built $f(h(\mathbf{x}, k))$ for the task, where h and f are the feature extractor and the task specific classifier, respectively, and $f(h(\mathbf{x}, k))$ is the output of the neural network for task k . We omit task-id k in f to simplify notation. In learning each task k , a task mask is also trained at the same time to protect the learned model of the task.

In testing for TIL, given the test sample \mathbf{x} with task-id k provided, CLOM uses the model for task k to classify \mathbf{x} ,

$$\hat{y} = \arg \max f(h(\mathbf{x}, k)) \quad (1)$$

For CIL (no task-id for each test sample \mathbf{x}), CLOM uses the model of every task and predicts the class \hat{y} using

$$\hat{y} = \arg \max \bigoplus_{1 \leq k \leq t} f(h(\mathbf{x}, k)), \quad (2)$$

where \bigoplus is the concatenation over the output space and t is the last task that has been learned. Eq. (2) essentially chooses the class with the highest classification score among the classes of all tasks learned so far. This works because the OOD detection model for a task will give very low score for a test sample that does not belong to the task. An overview of the prediction is illustrated in Fig. 1(a).

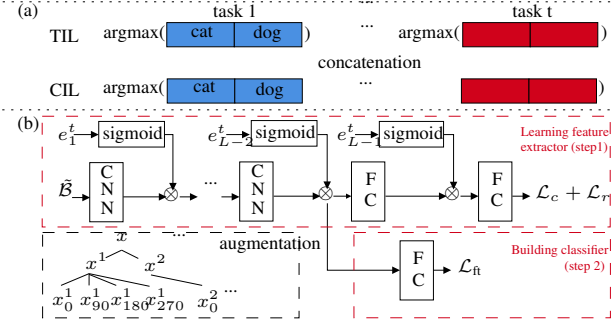


Figure 1. (a) Overview of TIL and CIL prediction. For TIL prediction, we consider only the output heads of the given task. For CIL prediction, we obtain all the outputs from task 1 to current task t , and choose the label over the concatenated output. (b) Training overview for a generic backbone. Each CNN includes a convolution layer with batch normalization, activation, or pooling depending on the configuration. In our experiments, we use AlexNet [35] and ResNet-18 [22]. We feed-forward augmented batch $\tilde{\mathcal{B}}$ consisting of rotated images of different views (x_0^1, x_0^2, \dots). We first train feature extractor by using contrastive loss (step 1). At each layer, binary mask is multiplied to the output of each convolution layer to learn important parameters of current task t . After training feature extractor, we fine-tune the OOD classifier (step 2).

Note that in fact any CIL model can be used as a TIL model if task-id is provided. The conversion is made by selecting the corresponding task classification heads in testing. However, a conversion from a TIL method to a CIL method is not obvious as the system requires task-id in testing. Some attempts have been made to predict the task-id in order to make a TIL method applicable to CIL problems. iTAML [52] requires the test samples to come in batches and each batch must be from a single task. This may not be practical as test samples usually come one by one. CCG [1] builds a separate network to predict the task-id, which is also prone to forgetting. Expert Gate [3] constructs a separate autoencoder for each task. Our CLOM classifies one test sample at a time and does not need to construct another network for task-id prediction.

In the following subsections, we present (1) how to build an OOD model for a task and use it to make a prediction, and (2) how to learn task masks to protect each model. An overview of the process is illustrated in Fig. 1(b).

3.1. Learning Each Task as OOD Detection

We borrow the latest OOD ideas based on contrastive learning [11, 21] and data augmentation due to their excellent performance [65]. Since this section focuses on how to learn a single task based on OOD detection, we omit the task-id unless necessary. The OOD training process is similar to that of contrastive learning. It consists of two steps: 1) learning the feature representation by the composite $g \circ h$, where h is a feature extractor and g is a projection to con-

trastive representation, and 2) learning a linear classifier f mapping the feature representation of h to the label space. In the following, we describe the training process: contrastive learning for feature representation learning (1), and OOD classifier building (2). We then explain how to make a prediction based on an ensemble method for both TIL and CIL settings, and how to further improve prediction using some saved past data.

3.1.1 Contrastive Loss for Feature Learning.

This is step 1 in Fig. 1(b). Supervised contrastive learning is used to try to repel data of different classes and align data of the same class more closely to make it easier to classify them. A key operation is data augmentation via transformations.

Given a batch of N samples, each sample x is first duplicated and each version then goes through *three initial augmentations* (also see Data Augmentation in Sec. 4.1) to generate two different views x^1 and x^2 (they keep the same class label as x). Denote the augmented batch by $\tilde{\mathcal{B}}$, which now has $2N$ samples. In [23, 65], it was shown that using image rotations is effective in learning OOD detection models because such rotations can effectively serve as out-of-distribution (OOD) training data. For each augmented sample $x \in \tilde{\mathcal{B}}$ with class y of a task, we rotate x by 90° , 180° , 270° to create three images, which are assigned *three new classes* y_1, y_2 , and y_3 , respectively. This results in a larger augmented batch $\tilde{\mathcal{B}}$. Since we generate three new images from each x , the size of $\tilde{\mathcal{B}}$ is $8N$. For each original class, we now have 4 classes. For a sample $x \in \tilde{\mathcal{B}}$, let $\tilde{\mathcal{B}}(x) = \tilde{\mathcal{B}} \setminus \{x\}$ and let $P(x) \subset \tilde{\mathcal{B}} \setminus \{x\}$ be a set consisting of the data of the same class as x distinct from x . The contrastive representation of a sample x is $z_x = g(h(x, t)) / \|g(h(x, t))\|$, where t is the current task. In learning, we minimize the supervised contrastive loss [31] of task t .

$$\mathcal{L}_c = \frac{1}{8N} \sum_{x \in \tilde{\mathcal{B}}} \frac{-1}{|P(x)|} \times \sum_{p \in P(x)} \log \frac{\exp(z_x \cdot z_p / \tau)}{\sum_{x' \in \tilde{\mathcal{B}}(x)} \exp(z_x \cdot z_{x'} / \tau)} \quad (3)$$

where τ is a scalar temperature, \cdot is dot product, and \times is multiplication. The loss is reduced by repelling z of different classes and aligning z of the same class more closely. \mathcal{L}_c basically trains a feature extractor with good representations for learning an OOD classifier.

3.1.2 Learning the Classifier.

This is step 2 in Fig. 1(b). Given the feature extractor h trained with the loss in Eq. (3), we *freeze* h and only *fine-tune* the linear classifier f , which is trained to predict the

classes of task t and the augmented rotation classes. f maps the feature representation to the label space in $\mathcal{R}^{4|\mathcal{C}^t|}$, where 4 is the number of rotation classes including the original data with 0° rotation and $|\mathcal{C}^t|$ is the number of original classes in task t . We minimize the cross-entropy loss,

$$\mathcal{L}_{\text{ft}} = -\frac{1}{|\tilde{\mathcal{B}}|} \sum_{(\mathbf{x}, y) \in \tilde{\mathcal{B}}} \log \tilde{p}(y|\mathbf{x}, t) \quad (4)$$

where ft indicates fine-tune, and

$$\tilde{p}(y|\mathbf{x}, t) = \text{softmax}(f(h(\mathbf{x}, t))) \quad (5)$$

where $f(h(\mathbf{x}, t)) \in \mathcal{R}^{4|\mathcal{C}^t|}$. The output $f(h(\mathbf{x}, t))$ includes the rotation classes. The linear classifier is trained to predict the original and the rotation classes.

3.1.3 Ensemble Class Prediction.

We describe how to predict a label $y \in \mathcal{C}^t$ (TIL) and $y \in \mathcal{C}$ (CIL) (\mathcal{C} is the set of original classes of all tasks) We assume all tasks have been learned and their models are protected by masks, which we discuss in the next subsection.

We discuss the prediction of class label y for a test sample \mathbf{x} in the TIL setting first. Note that the network $f \circ h$ in Eq. (5) returns logits for rotation classes (including the original task classes). Note also for each original class label $j_k \in \mathcal{C}^k$ (original classes) of a task k , we created three additional rotation classes. For class j_k , the classifier f will produce four output values from its four rotation class logits, i.e., $f_{j_k, 0}(h(\mathbf{x}_0, k))$, $f_{j_k, 90}(h(\mathbf{x}_{90}, k))$, $f_{j_k, 180}(h(\mathbf{x}_{180}, k))$, and $f_{j_k, 270}(h(\mathbf{x}_{270}, k))$, where 0, 90, 180, and 270 represent 0° , 90° , 180° , and 270° rotations respectively and \mathbf{x}_0 is the original \mathbf{x} . We compute an ensemble output $f_{j_k}(h(\mathbf{x}, k))$ for each class $j_k \in \mathcal{C}^k$ of task k ,

$$f_{j_k}(h(\mathbf{x}, k)) = \frac{1}{4} \sum_{\text{deg}} f_{j_k, \text{deg}}(h(\mathbf{x}_{\text{deg}}, k)). \quad (6)$$

The final TIL class prediction is made as follows (note, in TIL, task-id k is provided in testing),

$$\hat{y} = \arg \max_{j_k \in \mathcal{C}^k} \left\{ f_{j_k}(h(\mathbf{x}, k)) \right\} \quad (7)$$

We use Eq. (2) to make the CIL class prediction, where the final format $f(h(\mathbf{x}, k))$ for task k is the following vector:

$$f(h(\mathbf{x}, k)) = [f_1(h(\mathbf{x}, k), \dots, f_{|\mathcal{C}^k|}(h(\mathbf{x}, k))] \quad (8)$$

Our method so far memorizes no training samples and it already outperforms baselines (see Sec. 4.2).

3.1.4 Output Calibration with Memory.

The proposed method can make incorrect CIL prediction even with a perfect OOD model (rejecting every test sample that does not belong any class of the task). This happens because the task models are trained independently, and the outputs of different tasks may have different magnitudes. We use output calibration to ensure that the outputs are of similar magnitudes to be comparable by using some saved examples in a memory \mathcal{M} with limited budget. At each task k , we store a fraction of the validation data $\{(\mathbf{x}, y)\}$ into \mathcal{M}_k for output calibration and update the memory $\mathcal{M} \leftarrow \text{update}(\mathcal{M}, \mathcal{M}_k)$ by maintaining an equal number of samples per class. We will detail the memory budget in Sec. 4.3. Basically, we save the same number as the existing replay-based TIL/CIL methods.

After training the network f for the current task t , we freeze the model and use the saved data in \mathcal{M} to find the scaling and shifting parameters $(\sigma, \mu) \in \mathcal{R}^t \times \mathcal{R}^t$ to calibrate the after-ensemble classification output $f(h(\mathbf{x}, k))$ (Eq. (8)) (i.e., using $\sigma_k f(h(\mathbf{x}, k)) + \mu_k$) for each task k by minimizing the cross-entropy loss,

$$\mathcal{L}_{\text{cal}} = -\frac{1}{|\mathcal{M}|} \sum_{(\mathbf{x}, y) \in \mathcal{M}} \log p(y|\mathbf{x}) \quad (9)$$

where $p(c|\mathbf{x})$ is computed using softmax,

$$\text{softmax} \left[\sigma_1 f(h(\mathbf{x}, 1)) + \mu_1; \dots; \sigma_t f(h(\mathbf{x}, t)) + \mu_t \right] \quad (10)$$

Clearly, parameters σ_k and μ_k do not change classification within task k (TIL), but calibrates the outputs such that the ensemble outputs from all tasks are in comparable magnitudes. For CIL inference at the test time, we make prediction by (following Eq. (2)),

$$\hat{y} = \arg \max_{1 \leq k \leq t} \bigoplus [\sigma_k f(h(\mathbf{x}, k)) + \mu_k] \quad (11)$$

3.2. Protecting OOD Model of Each Task Using Masks

We now discuss the task mask mechanism based on hard attentions [61] for protecting the OOD model of each task to deal with CF. In learning the OOD model for each task, CLOM at the same time also trains a *mask* or *hard attention* for each layer. To protect the shared feature extractor from previous tasks, their masks are used to block those important neurons so that the new task learning will not interfere with the parameters learned for previous tasks.

The main idea is to use sigmoid to approximate a 0-1 step function as *hard attention* to mask (or block) or unmask (unblock) the information flow to protect the parameters learned for each previous task.

The hard attention (mask) at layer l and task t is defined as

$$\mathbf{a}_l^t = u(s\mathbf{e}_l^t) \quad (12)$$

where u is the sigmoid function, s is a scalar, and \mathbf{e}_l^t is a *learnable* embedding of the task-id input t . The attention is element-wise multiplied to the output \mathbf{h}_l of layer l as

$$\mathbf{h}_l' = \mathbf{a}_l^t \otimes \mathbf{h}_l \quad (13)$$

as depicted in Fig. 1(b). The sigmoid function u converges to a 0-1 step function as s goes to infinity. Since the true step function is not differentiable, a fairly large s is chosen to achieve a differentiable pseudo step function based on sigmoid (see Appendix 4 for choice of s). The pseudo binary value of the attention determines how much information can flow forward and backward between adjacent layers.

Denote $\mathbf{h}_l = \text{ReLU}(\mathbf{W}_l\mathbf{h}_{l-1} + \mathbf{b}_l)$, where ReLU is the rectifier function. For units (neurons) of attention \mathbf{a}_l^t with zero values, we can freely change the corresponding parameters in \mathbf{W}_l and \mathbf{b}_l without affecting the output \mathbf{h}_l' . For units of attention with non-zero values, changing the parameters will affect the output \mathbf{h}_l' for which we need to protect from gradient flow in backpropagation to prevent forgetting.

Specifically, during training the new task t , we update parameters according to the attention so that the important parameters for past tasks ($1, \dots, t-1$) are unmodified. Denote the accumulated attentions (masks) of all past tasks by

$$\mathbf{a}_l^{<t} = \max(\mathbf{a}_l^{<t-1}, \mathbf{a}_l^{t-1}) \quad (14)$$

where $\mathbf{a}_l^{<t}$ is the hard attentions of layer l of all previous tasks, and \max is an element-wise maximum¹ and \mathbf{a}_l^0 is a zero vector. Then the modified gradient is the following,

$$\nabla w'_{ij,l} = \left(1 - \min\left(a_{i,l}^{<t}, a_{j,l-1}^{<t}\right)\right) \nabla w_{ij,l} \quad (15)$$

where $a_{i,l}^{<t}$ indicates i 'th unit of $\mathbf{a}_l^{<t}$ and $l = 1, \dots, L-1$. This reduces the gradient if the corresponding units' attentions at layers l and $l-1$ are non-zero². We do not apply hard attention on the last layer L because it is a task-specific layer.

To encourage sparsity in \mathbf{a}_l^t and parameter sharing with $\mathbf{a}_l^{<t}$, a regularization (\mathcal{L}_r) for attention at task t is defined as

$$\mathcal{L}_r = \lambda_t \frac{\sum_l \sum_i a_{i,l}^t (1 - a_{i,l}^{<t})}{\sum_l \sum_i (1 - a_{i,l}^{<t})} \quad (16)$$

¹Some parameters from different tasks can be shared, which means some hard attention masks can be shared.

²By construction, if \mathbf{a}_l^t becomes $\mathbf{1}$ for all layers, the gradients are zero and the network is at maximum capacity. However, the network capacity can increase by adding more parameters.

where λ_t is a scalar hyperparameter. For flexibility, we denote λ_t for each task t . However, in practice, we use the same λ_t for all $t \neq 1$. The final objective to be minimized for task t with hard attention is (see Fig. 1(b))

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_r \quad (17)$$

where \mathcal{L}_c is the contrastive loss function (Eq. (3)). By protecting important parameters from changing during training, the neural network effectively alleviates CF.

4. Experiments

Evaluation Datasets: Four image classification CL benchmark datasets are used in our experiments.

(1) **MNIST**³: handwritten digits of 10 classes (digits) with 60,000 examples for training and 10,000 examples for testing.

(2) **CIFAR-10** [34]⁴: 60,000 32x32 color images of 10 classes with 50,000 for training and 10,000 for testing.

(3) **CIFAR-100** [34]⁵: 60,000 32x32 color images of 100 classes with 500 images per class for training and 100 per class for testing.

(4) **Tiny-ImageNet** [36]⁶: 120,000 64x64 color images of 200 classes with 500 images per class for training and 50 images per class for validation, and 50 images per class for testing. Since the test data has no labels in this dataset, we use the validation data as the test data as in [40].

Baseline Systems: We compare our CLOM with both the classic and the most recent state-of-the-art CIL and TIL methods. We also include CLOM(-c), which is CLOM without calibration (which already outperforms the baselines).

For CIL baselines, we consider seven *replay* methods, **LwFR** (replay version of LwF [38] with better results [42]), **iCaRL** [54], **BiC** [70], **A-RPS** [51], **Mnemonics** [42], **DER++** [5], and **Co²L** [8]; one *pseudo-replay* method **CCG** [1]⁷; one *orthogonal projection* method **OWM** [74]; and a *multi-classifier* method **MUC** [40]; and a *prototype augmentation method* **PASS** [77]. OWM, MUC, and PASS *do not save any samples* from previous tasks.

TIL baselines include **HAT** [61], **HyperNet** [67], and **SupSup** [68]. As noted earlier, CIL methods can also be used for TIL. In fact, TIL methods may be used for CIL too, but the results are very poor. We include them in our comparison

³<http://yann.lecun.com/exdb/mnist/>

⁴<https://www.cs.toronto.edu/~kriz/cifar.html>

⁵<https://www.cs.toronto.edu/~kriz/cifar.html>

⁶<http://tiny-imagenet.herokuapp.com>

⁷iTAML [52] is not included as they require a batch of test data from the same task to predict the task-id. When each batch has only one test sample, which is our setting, it is very weak. For example, iTAML TIL/CIL accuracy is only 35.2%/33.5% on CIFAR100 10 tasks. Expert Gate (EG) [3] is also weak. For example, its TIL/CIL accuracy is 87.2/43.2 on MNIST 5 tasks. Both iTAML and EG are much weaker than many baselines.

Method	MNIST-5T		CIFAR10-5T		CIFAR100-10T		CIFAR100-20T		T-ImageNet-5T		T-ImageNet-10T		Average	
	TIL	CIL	TIL	CIL	TIL	CIL	TIL	CIL	TIL	CIL	TIL	CIL	TIL	CIL
CIL Systems														
OWM	99.7±0.03	95.8±0.13	85.2±0.17	51.7±0.06	59.9±0.84	29.0±0.72	65.4±0.07	24.2±0.11	22.4±0.87	10.0±0.55	28.1±0.55	8.6±0.42	60.1	36.5
MUC	99.9±0.02	74.6±0.45	95.2±0.24	53.6±0.95	76.9±1.27	30.0±1.37	73.7±1.27	14.4±0.93	55.8±0.26	33.6±0.18	47.2±0.22	17.4±0.17	74.8	37.3
PASS†	99.5±0.14	76.6±1.67	83.8±0.68	47.3±0.97	72.4±1.23	36.8±1.64	76.9±0.77	25.3±0.81	50.3±1.97	28.9±1.36	47.6±0.38	18.7±0.58	71.8	38.9
LwFR	99.9 ±0.09	85.0±3.05	95.2±0.30	54.7±1.18	86.2±1.00	45.3±0.75	89.0±0.45	44.3±0.46	56.4±0.48	32.2±0.50	55.3±0.35	24.3±0.26	80.3	47.6
iCaRL*	99.9 ±0.09	96.0±0.42	94.9±0.34	63.4±1.11	84.2±1.04	51.4±0.99	85.7±0.68	47.8±0.48	54.3±0.59	37.0±0.41	52.7±0.37	28.3±0.18	78.6	54.0
Mnemonics†*	99.9 ±0.03	96.3±0.36	94.5±0.46	64.1±1.47	82.3±0.30	51.0±0.34	86.2±0.46	47.6±0.74	54.8±0.16	37.1±0.46	52.9±0.66	28.5±0.72	78.5	54.1
BiC	99.9 ±0.04	85.1±1.84	91.1±0.82	57.1±1.09	87.6±0.28	51.3±0.59	90.3±0.26	40.1±0.77	44.7±0.71	20.2±0.31	50.3±0.65	21.2±0.46	77.3	45.8
DER++	99.7±0.08	95.3±0.69	92.2±0.48	66.0±1.27	84.2±0.47	55.3±0.10	86.6±0.50	46.6±1.44	58.0±0.52	36.0±0.42	59.7±0.6	30.5±0.30	80.1	55.0
A-RPS						60.8								
CCG		97.3		70.1				53.5						
Co ² L			93.4	65.6										
TIL Systems														
HAT	99.9 ±0.02	81.9±3.73	96.7±0.18	62.7±1.46	84.0±0.23	41.1±0.93	85.0±0.85	26.0±0.83	61.2±0.72	38.5±1.85	63.8±0.41	29.8±0.65	81.8	46.6
HyperNet	99.7±0.05	49.1±5.52	94.9±0.54	47.4±5.78	77.3±0.45	29.7±2.19	83.0±0.60	19.4±1.44	23.8±1.21	8.8±0.98	27.8±0.86	5.8±0.56	67.8	26.7
SupSup	99.6±0.08	71.3±4.06	94.8±0.18	49.8±1.40	88.0±0.13	44.2±0.80	91.6±0.17	35.0±0.08	64.2±0.66	43.0±0.82	68.2±0.26	36.5±0.30	84.4	46.6
CLOM(-c)	99.9 ±0.00	94.4±0.26	98.7 ±0.06	87.8±0.71	92.0 ±0.37	63.3±1.00	94.3 ±0.06	54.6±0.92	68.4 ±0.16	45.7±0.26	72.4 ±0.21	47.1±0.18	87.6	65.5
CLOM	99.9 ±0.00	96.9±0.30	98.7 ±0.06	88.0 ±0.48	92.0 ±0.37	65.2 ±0.71	94.3 ±0.06	58.0 ±0.45	68.4 ±0.16	51.7 ±0.37	72.4 ±0.21	47.6 ±0.32	87.6	67.9

Table 1. Average accuracy over all classes after the last task is learned. -xT: x number of tasks. †: In their original paper, PASS and Mnemonics use the first half of classes to pre-train before CL. Their results are 50.1% and 53.5% on CIFAR100-10T respectively, but they are still lower than CLOM without pre-training. In our experiments, no pre-training is used for fairness. *: **iCaRL** and **Mnemonics** give both the final average accuracy as here and the *average incremental accuracy* in the original papers. We report the *average incremental accuracy* and *network size* in Appendix 1 and 2, respectively. The last two columns show the average TIL and CIL accuracy of each method over all datasets.

4.1. Training Details

For all experiments, we use 10% of training data as the validation set to grid-search for good hyperparameters. For minimizing the contrastive loss, we use LARS [73] for 700 epochs with initial learning rate 0.1. We linearly increase the learning rate by 0.1 per epoch for the first 10 epochs until 1.0 and then decay it by cosine scheduler [44] after 10 epochs without restart as in [11, 65]. For fine-tuning the classifier f , we use SGD for 100 epochs with learning rate 0.1 and reduce the learning rate by 0.1 at 60, 75, and 90 epochs. The full set of hyperparameters is given in Appendix 4.

We follow the recent baselines (A-RPS, DER++, PASS and Co²L) and use the same class split and backbone architecture for both CLOM and baselines.

For MNIST and CIFAR-10, we split 10 classes into 5 tasks where each task has 2 classes in consecutive order. We save 20 random samples per class from the validation set for output calibration. This number is commonly used in replay methods [42, 51, 54, 70]. MNIST consists of single channel images of size 1x28x28. Since the contrastive learning [11] relies on color changes, we copy the channel to make 3-channels. For MNIST and CIFAR-10, we use AlexNet-like architecture [35] and ResNet-18 [22] respectively for both CLOM and baselines.

For CIFAR-100, we conduct two experiments. We split 100 classes into 10 tasks and 20 tasks where each task has 10 and 5 classes, respectively, in consecutive order. We use 2000 memory budget as in [54], saving 20 random samples per class from the validation set for output calibration. We

use the same ResNet-18 structure for CLOM and baselines, but we increase the number of channels twice to learn more tasks.

For Tiny-ImageNet, we follow [40] and resize the original images of size 3x64x64 to 3x32x32 so that the same ResNet-18 of CIFAR-100 experiment setting can be used. We split 200 classes into 5 tasks (40 classes per task) and 10 tasks (20 classes per task) in consecutive order, respectively. To have the same memory budget of 2000 as for CIFAR-100, we save 10 random samples per class from the validation set for output calibration.

Data Augmentation. For baselines, we use data augmentations used in their original papers. For CLOM, following [11, 65], we use *three initial augmentations* (see Sec. 3.1) (i.e., *horizontal flip*, *color change (color jitter and grayscale)*, and *Inception crop* [64]) and four *rotations* (see Sec. 3.1). Specific details about these transformations are given in Appendix 3.

4.2. Results and Comparative Analysis

As in existing works, we evaluate each method by two metrics: *average classification accuracy* on all classes after training the last task, and *average forgetting rate* [42], $F^t = \frac{1}{t-1} \sum_{j=1}^{t-1} A_j^{\text{init}} - A_j^t$, where A_j^{init} is the j 'th task's accuracy of the network right after the j 'th task is learned and A_j^t is the accuracy of the network on the j 'th task data after learning the last task t . We report the forgetting rate after the final task t . Our results are averages of 5 random runs.

We present the main experiment results in Tab. 1. The last

two columns give the average TIL/CIL results of each system/row. For A-RPS, CCG, and Co²L, we copy the results from their original papers as their codes are not released to the public or the public code cannot run on our system. The rows are grouped by CIL and TIL methods.

CIL Results Comparison. Tab. 1 shows that CLOM and its memory free method CLOM(-c) achieve much higher CIL accuracy except for MNIST for which CLOM is slightly weaker than CCG by 0.4%, but CLOM’s result on CIFAR10-5T is about 18% greater than CCG. For other datasets, CLOM improves by similar margins. This is in contrast to the baseline TIL systems that are incompetent at the CIL setting when classes are predicted using Eq. (2). Even *without* calibration, CLOM(-c) already outperforms all the baselines by large margins.

TIL Results Comparison. The gains by CLOM and CLOM(-c) over the baselines are also great in the TIL setting. CLOM and CLOM(-c) are the same as the output calibration does not affect TIL performance. For the two large datasets CIFAR100 and T-ImageNet, CLOM gains by large margins. This is due to contrastive learning and the OOD model. The replay based CIL methods (LwF.R, iCaRL, Mnemonics, BiC, and DER++) perform reasonably in the TIL setting, but our CLOM and CLOM(-c) are much better due to task masks which can protect previous models better with little CF.




Figure 2. Average forgetting rate (%) in the TIL setting as CLOM is a TIL system. The lower the value, the better the method is. CIL/TIL systems are shaded in blue/red, respectively (*best viewed in color*). A negative value indicates the task accuracy has increased from the initial accuracy.

Comparison of Forgetting Rate. Fig. 2 shows the average forgetting rate of each method in the TIL setting. The CIL systems suffer from more forgetting as they are not designed for the TIL setting, which results in lower TIL accuracy (Tab. 1). The TIL systems are highly effective at preserving previous within-task knowledge. This results in higher TIL accuracy on large dataset such as T-ImageNet, but they collapse when task-id is not provided (the CIL setting) as shown in Tab. 1. CLOM is robust to forgetting as a TIL system and it also functions well without task-id.

We report only the forgetting rate in the TIL setting be-

	CIFAR10-5T				CIFAR100-10T			
	AUC	TaskDR	TIL	CIL	AUC	TaskDR	TIL	CIL
SupSup	78.9	26.2	95.3	26.2	76.7	34.3	85.2	33.1
SupSup (OOD in CLOM)	88.9	82.3	97.2	81.5	84.9	63.7	90.0	62.1
CLOM (ODIN)	82.9	63.3	96.7	62.9	77.9	43.0	84.0	41.3
CLOM	92.2	88.5	98.7	88.0	85.0	66.8	92.0	65.2
CLOM (w/o OOD)	90.3	83.9	98.1	83.3	82.6	59.5	89.8	57.5

Table 2. TIL and CIL results improve with better OOD detection. Column AUC gives the average AUC score for the OOD detection method as used within each system on the left. Column TaskDR gives *task detection rate*. TIL and CIL results are average accuracy values. SupSup and CLOM variants are calibrated with 20 samples per class.

cause our CLOM is essentially a TIL method and not a CIL system by design. The degrading CIL accuracy of CLOM is mainly because the OOD model for each task is not perfect.

4.3. Ablation Studies

Better OOD for better continual learning. We show that (1) an existing CL model can be improved by a good OOD model and (2) CLOM’s results will deteriorate if a weaker OOD model is applied. To isolate effect of OOD detection on changes in CIL performance, we further define *task detection* and *task detection rate*. For a test sample from a class of task j , if it is predicted to a class of task m and $j = m$, the task detection is correct for this test instance. The *task detection rate* $\sum_{\mathbf{x} \in \mathcal{D}^{\text{test}}} 1_{j=m} / N$, where N is the total number of test instances in $\mathcal{D}^{\text{test}}$, is the rate of correct task detection. We measure the performance of OOD detection using AUC (Area Under the ROC Curve) averaged over all tasks. AUC is the main measure used in OOD detection papers. We conduct experiments on CIFAR10-5T and CIFAR100-10T.

For (1), we use the TIL baseline SupSup as it displays a strong TIL performance and is robust to forgetting like CLOM. We replace SupSup’s task learner with the OOD model in CLOM. Tab. 2 shows that the OOD method in CLOM improves SupSup (SupSup (OOD in CLOM)) greatly. It shows that our approach is applicable to different TIL systems.

For (2), we replace CLOM’s OOD method with a weaker OOD method ODIN [39]. We see in Tab. 2 that task detection rate, TIL, and CIL results all drop markedly with ODIN (CLOM (ODIN)).

CLOM without OOD detection. In this case, CLOM uses contrastive learning and data augmentation, but does not use the rotation classes in classification. Note that the rotation classes are basically regarded as OOD data in training and for OOD detection in testing. CLOM (w/o OOD) in Tab. 2 represents this CLOM variant. We see that CLOM (w/o OOD) is much weaker than the full CLOM. This indicates that the improved results of CLOM over baselines are not only due to contrastive learning and data augmentation

$ \mathcal{M} $	(a)	(b)	s	F^5	AUC	CIL
0	63.3±1.00	54.6±0.92	1	48.6	58.8	10.0
5	64.9±0.67	57.7±0.50	100	13.3	82.7	67.7
10	65.0±0.71	57.8±0.53	300	8.2	83.3	72.0
15	65.1±0.71	57.9±0.44	500	0.2	91.8	87.2
20	65.2±0.71	58.0±0.45	700	0.1	92.2	88.0

Table 3. (Left) shows changes of accuracy with the number of samples saved per class for output calibration. (a) and (b) are CIFAR100-10T and CIFAR100-20T, respectively. $|\mathcal{M}| = k$ indicates k samples are saved per class. (Right) shows that weaker forgetting avoidance mechanism results in larger forgetting and lower AUC, thus lower CIL. For $s = 1$, the pseudo-step function becomes the standard sigmoid, thus parameters are hardly protected. F^5 is the forgetting rate over 5 tasks.

Aug.	CIFAR10-5T		CIFAR100-10T	
	TIL	CIL	TIL	CIL
Hflip	(93.1, 95.3)	(49.1, 72.7)	(77.6, 84.0)	(31.1, 47.0)
Color	(91.7, 94.6)	(50.9, 70.2)	(67.2, 77.4)	(28.7, 41.8)
Crop	(96.1, 97.3)	(58.4, 79.4)	(84.1, 89.3)	(41.6, 60.3)
All	(97.6, 98.7)	(74.0, 88.0)	(88.1, 92.0)	(50.2, 65.2)

Table 4. Accuracy of CLOM variants when a single or all augmentations are applied. Hflip: horizontal flip; Color: color jitter and grayscale; Crop: Inception [64]. (num1, num2): accuracy without and with rotation.

but also significantly due to OOD detection.

Effect of the number of saved samples for calibration.

Tab. 3 (left) reveals that the output calibration is still effective with a small number of saved samples per class ($|\mathcal{M}|$). For both CIFAR100-10T and CIFAR100-20T, CLOM achieves competitive performance by using only 5 samples per class. The accuracy improves and the variance decreases with the number of saved samples.

Effect of s in Eq. (12) on forgetting of CLOM. We need to use a strong forgetting avoidance mechanism for CLOM to be functional. Using CIFAR10-5T, we show how CLOM performs with different values of s in hard attention or masking. The larger s value, the stronger protection is used. Tab. 3 (right) shows that average AUC and CIL decrease as the forgetting rate increases. This also supports the result in Tab. 2 that SupSup improves greatly with the OOD method in CLOM as it is also robust to forgetting. PASS and Co²L underperform despite they also use rotation or contrastive loss as their forgetting avoidance mechanisms are weak.

Effect of data augmentations. For data augmentation, we use three initial augmentations (*i.e.*, *horizontal flip*, *color change* (color jitter and grayscale), *Inception crop* [64]), which are commonly used in contrastive learning to build a single model. We additionally use rotation for OOD data in training. To evaluate the contribution of each augmentation

when task models are trained sequentially, we train CLOM using one augmentation. We do not report their effects on forgetting as we experience rarely any forgetting (Fig. 2 and Tab. 3). Tab. 4 shows that the performance is lower when only a single augmentation is applied. When all augmentations are applied, the TIL/CIL accuracies are higher. The rotation always improves the result when it is combined with other augmentations. More importantly, when we use crop and rotation, we achieve higher CIL accuracy (79.4/60.3% for CIFAR10-5T/CIFAR100-10T) than we use all augmentations without rotation (74.0/50.2%). This shows the efficacy of rotation in our system.

5. Conclusions

This paper proposed a novel continual learning (CL) method called CLOM based on OOD detection and task masking that can perform both task incremental learning (TIL) and class incremental learning (CIL). Regardless whether it is used for TIL or CIL in testing, the training process is the same, which is an advantage over existing CL systems as they focus on either CIL or TIL and have limitations on the other problem. Experimental results showed that CLOM outperforms both state-of-the-art TIL and CIL methods by very large margins. In our future work, we will study ways to improve efficiency and also accuracy.

Limitation of our work: One limitation of this work is that since CLOM needs to build a good OOD model, it is unclear whether it can be effectively applied to the online continual learning setting where the data comes in a stream and can be visited once. Incrementally building an accurate OOD model can be challenging as the current method needs to train in a large number of epochs. This is a direction that we plan to explore in our future work.

Potential negative societal impact: We do not see any except that like many other machine learning algorithms, our proposed method could be affected by bias in the input data as CLOM does not deal with fairness or bias in the data.

References

- [1] Davide Abati, Jakub Tomczak, Tijmen Blankevoort, Simone Calderara, Rita Cucchiara, and Ehteshami Bejnordi. Conditional channel gated networks for task-aware continual learning. In *CVPR*, pages 3931–3940, 2020. 3, 5
- [2] Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. Uncertainty-based continual learning with adaptive regularization. In *NeurIPS*, 2019. 2
- [3] Rahaf Aljundi, Punarjay Chakravarty, and Tinne Tuytelaars. Expert gate: Lifelong learning with a network of experts. In *CVPR*, 2017. 2, 3, 5
- [4] Saikiran Bulusu, Bhavya Kailkhura, Bo Li, Pramod K Varshney, and Dawn Song. Anomalous instance detection in deep learning: A survey. *arXiv preprint arXiv:2003.06979*, 2020. 1, 2

- [5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and SIMONE CALDERARA. Dark experience for general continual learning: a strong, simple baseline. In *NeurIPS*, 2020. 2, 5
- [6] Raffaello Camoriano, Giulia Pasquale, Carlo Ciliberto, Lorenzo Natale, Lorenzo Rosasco, and Giorgio Metta. Incremental robot learning of new objects with fixed update time. In *ICRA*, 2017. 2
- [7] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, pages 233–248, 2018. 2
- [8] Hyuntak Cha, Jaeho Lee, and Jinwoo Shin. Co2l: Contrastive continual learning. In *ICCV*, 2021. 2, 5
- [9] Arslan Chaudhry, Naeemullah Khan, Puneet K. Dokania, and Philip H. S. Torr. Continual learning in low-rank orthogonal subspaces, 2020. 2
- [10] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *ICLR*, 2019. 2
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 1, 2, 3, 6
- [12] Zhiyuan Chen and Bing Liu. *Lifelong Machine Learning*. Morgan & Claypool Publishers, 2018. 1
- [13] Cyprien de Masson d’Autume, Sebastian Ruder, Lingpeng Kong, and Dani Yogatama. Episodic memory in lifelong language learning. In *NeurIPS*, 2019. 2
- [14] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyang Wu, and Rama Chellappa. Learning without memorizing. In *CVPR*, 2019. 1, 2
- [15] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. 2
- [16] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 2
- [17] Alexander Gepperth and Cem Karaoguz. A bio-inspired incremental learning architecture for applied perceptual problems. *Cognitive Computation*, 8(5):924–934, 2016. 2
- [18] Izhak Golan and Ran El-Yaniv. Deep anomaly detection using geometric transformations. In *NeurIPS*, 2018. 1
- [19] Yiduo Guo, Wenpeng Hu, Dongyan Zhao, and Bing Liu. Adaptive orthogonal projection for batch and online continual learning. In *Proceedings of AAAI-2022*, 2022. 2
- [20] Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. *arXiv preprint arXiv:1910.02509*, 2019. 2
- [21] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3, 6
- [23] Dan Hendrycks, Mantas Mazeika, Saurav Kadavath, and Dawn Song. Using self-supervised learning can improve model robustness and uncertainty. In *NeurIPS*, pages 15663–15674, 2019. 1, 3
- [24] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. 2
- [25] Wenpeng Hu, Zhou Lin, Bing Liu, Chongyang Tao, Zhengwei Tao, Jinwen Ma, Dongyan Zhao, and Rui Yan. Overcoming catastrophic forgetting for continual learning via model adaptation. In *ICLR*, 2019. 2
- [26] Ching-Yi Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In *NeurIPS*, volume 32, 2019. 2
- [27] Heechul Jung, Jeongwoo Ju, Minju Jung, and Junmo Kim. Less-forgetting learning in deep neural networks. *arXiv preprint arXiv:1607.00122*, 2016. 2
- [28] Nitin Kamra, Umang Gupta, and Yan Liu. Deep Generative Dual Memory Network for Continual Learning. *arXiv preprint arXiv:1710.10368*, 2017. 2
- [29] Zixuan Ke, Bing Liu, and Xingchang Huang. Continual learning of a mixed sequence of similar and dissimilar tasks. In *NeurIPS*, 2020. 2
- [30] Ronald Kemker and Christopher Kanan. FearNet: Brain-Inspired Model for Incremental Learning. In *ICLR*, 2018. 2
- [31] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362*, 2020. 3
- [32] Gyuhak Kim and Bing Liu. Continual learning via principal components projection, 2020. 2
- [33] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, and Others. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. 2
- [34] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. *Technical Report TR-2009*, University of Toronto, Toronto., 2009. 5
- [35] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 3, 6
- [36] Y. Le and X. Yang. Tiny imagenet visual recognition challenge, 2015. 5
- [37] Kibok Lee, Kimin Lee, Jinwoo Shin, and Honglak Lee. Overcoming catastrophic forgetting with unlabeled data in the wild. In *CVPR*, 2019. 2
- [38] Zhizhong Li and Derek Hoiem. Learning Without Forgetting. In *ECCV*, pages 614–629. Springer, 2016. 2, 5
- [39] Shiyu Liang, Yixuan Li, and R. Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018. 7
- [40] Yu Liu, Sarah Parisot, Gregory Slabaugh, Xu Jia, Ales Leonardis, and Tinne Tuytelaars. More classifiers, less forgetting: A generic multi-classifier paradigm for incremental

- learning. In *ECCV*, pages 699–716. Springer International Publishing, 2020. [2](#), [5](#), [6](#)
- [41] Yaoyao Liu, Bernt Schiele, and Qianru Sun. Adaptive aggregation networks for class-incremental learning. In *CVPR*, 2021. [2](#)
- [42] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *CVPR*, 2020. [5](#), [6](#)
- [43] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. In *NeurIPS*, pages 6470–6479, 2017. [2](#)
- [44] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [6](#)
- [45] Arun Mallya and Svetlana Lazebnik. PackNet: Adding Multiple Tasks to a Single Network by Iterative Pruning. *arXiv preprint arXiv:1711.05769*, 2017. [2](#)
- [46] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. [1](#)
- [47] Fei Mi, Lingjing Kong, Tao Lin, Kaicheng Yu, and Boi Faltings. Generalized class incremental learning. In *CVPR*, 2020. [2](#)
- [48] Aaron Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018. [1](#)
- [49] Oleksiy Ostapenko, Mihai Puscas, Tassilo Klein, Patrick Jahnichen, and Moin Nabi. Learning to remember: A synaptic plasticity driven framework for continual learning. In *CVPR*, pages 11321–11329, 2019. [2](#)
- [50] German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019. [1](#)
- [51] Jathushan Rajasegaran, Munawar Hayat, Salman Khan, Fahad Shahbaz Khan, Ling Shao, and Ming-Hsuan Yang. An adaptive random path selection approach for incremental learning, 2020. [2](#), [5](#), [6](#)
- [52] Jathushan Rajasegaran, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Mubarak Shah. itaml: An incremental task-agnostic meta-learning approach. In *CVPR*, 2020. [3](#), [5](#)
- [53] Amal Rannen Ep Triki, Rahaf Aljundi, Matthew Blaschko, and Tinne Tuytelaars. Encoder based lifelong learning. In *ICCV*, 2017. [2](#)
- [54] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, and Christoph H Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, pages 5533–5542, 2017. [2](#), [5](#), [6](#)
- [55] Hippolyt Ritter, Aleksandar Botev, and David Barber. Online structured laplace approximations for overcoming catastrophic forgetting. In *NeurIPS*, 2018. [2](#)
- [56] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Greg Wayne. Experience replay for continual learning. In *NeurIPS*, 2019. [2](#)
- [57] Mohammad Rostami, Soheil Kolouri, and Praveen K. Pilly. Complementary learning for overcoming catastrophic forgetting using experience replay. In *IJCAI*, 2019. [2](#)
- [58] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. [2](#)
- [59] Jonathan Schwarz, Jelena Luketina, Wojciech M Czarnecki, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. *arXiv preprint arXiv:1805.06370*, 2018. [2](#)
- [60] Ari Seff, Alex Beatson, Daniel Suo, and Han Liu. Continual learning in generative adversarial nets. *arXiv preprint arXiv:1705.08395*, 2017. [2](#)
- [61] Joan Serra, Dídac Surís, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *ICML*, 2018. [1](#), [2](#), [4](#), [5](#)
- [62] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NIPS*, pages 2994–3003, 2017. [2](#)
- [63] Pravendra Singh, Vinay Kumar Verma, Pratik Mazumder, Lawrence Carin, and Piyush Rai. Calibrating cnns for lifelong learning. *NeurIPS*, 2020. [2](#)
- [64] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. [6](#), [8](#)
- [65] Jihoon Tack, Sangwoo Mo, Jongheon Jeong, and Jinwoo Shin. Csi: Novelty detection via contrastive learning on distributionally shifted instances. In *NeurIPS*, 2020. [1](#), [2](#), [3](#), [6](#)
- [66] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. [1](#)
- [67] Johannes von Oswald, Christian Henning, João Sacramento, and Benjamin F Grewe. Continual learning with hypernetworks. *ICLR*, 2020. [2](#), [5](#)
- [68] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Anirudha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *NeurIPS*, volume 33, pages 15173–15184. Curran Associates, Inc., 2020. [2](#), [5](#)
- [69] Chenshen Wu, Luis Herranz, Xialei Liu, Joost van de Weijer, Bogdan Raducanu, et al. Memory replay gans: Learning to generate new categories without forgetting. In *NeurIPS*, 2018. [2](#)
- [70] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *CVPR*, 2019. [2](#), [5](#), [6](#)
- [71] Ju Xu and Zhanxing Zhu. Reinforced continual learning. In *NeurIPS*, 2018. [2](#)
- [72] Jaehong Yoon, Saehoon Kim, Eunho Yang, and Sung Ju Hwang. Scalable and order-robust continual learning with additive parameter decomposition. In *ICLR*, 2020. [2](#)
- [73] Yang You, Igor Gitman, and Boris Ginsburg. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888*, 2017. [6](#)
- [74] Guanxiong Zeng, Yang Chen, Bo Cui, and Shan Yu. Continuous learning of context-dependent processing in neural networks. *Nature Machine Intelligence*, 2019. [2](#), [5](#)

- [75] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *ICML*, pages 3987–3995, 2017. [2](#)
- [76] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shu-Tao Xia. Maintaining discrimination and fairness in class incremental learning. In *CVPR*, pages 13208–13217, 2020. [2](#)
- [77] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *CVPR*, 2021. [2](#), [5](#)