

# Few-Shot Lifelong Learning

Pratik Mazumder\*<sup>1</sup>, Pravendra Singh\*<sup>2</sup>, Piyush Rai<sup>1</sup>

<sup>1</sup> Department of Computer Science and Engineering, IIT Kanpur, India

<sup>2</sup> Independent Researcher, India

pratikm@cse.iitk.ac.in, pravendra1988@gmail.com, piyush@cse.iitk.ac.in

## Abstract

Many real-world classification problems often have classes with very few labeled training samples. Moreover, all possible classes may not be initially available for training, and may be given incrementally. Deep learning models need to deal with this two-fold problem in order to perform well in real-life situations. In this paper, we propose a novel Few-Shot Lifelong Learning (FSL) method that enables deep learning models to perform lifelong/continual learning on few-shot data. Our method selects very few parameters from the model for training every new set of classes instead of training the full model. This helps in preventing overfitting. We choose the few parameters from the model in such a way that only the currently unimportant parameters get selected. By keeping the important parameters in the model intact, our approach minimizes catastrophic forgetting. Furthermore, we minimize the cosine similarity between the new and the old class prototypes in order to maximize their separation, thereby improving the classification performance. We also show that integrating our method with self-supervision improves the model performance significantly. We experimentally show that our method significantly outperforms existing methods on the miniImageNet, CIFAR-100, and CUB-200 datasets. Specifically, we outperform the state-of-the-art method by an absolute margin of 19.27% for the CUB dataset.

## Introduction

Deep learning models have successfully matched human beings in many real-world problems. As a result, the number and diversity of applications of deep learning are increasing at a rapid rate. However, deep learning models require training on a large amount of labeled data. Labeled data is not always available for many real-world problems, and manually labeling data is a costly and time-consuming process. Therefore, recent works have investigated few-shot learning methods (Snell, Swersky, and Zemel 2017; Sung et al. 2018; Finn, Abbeel, and Levine 2017), which involve a specialized training of models to help them perform well even for classes with very few training samples.

Another common characteristic of real-world problems is that all the training data may not be available initially (Rebuffi et al. 2017; Li and Hoiem 2018; Castro et al.

2018). New sets of classes may become available incrementally. Therefore, the model has to perform lifelong/continual learning in order to perform well in such settings. The lifelong learning problem generally involves training a model on a sequence of disjoint sets of classes (task) and learning a joint classifier for all the encountered classes. This setting is also known as the class-incremental setting of lifelong learning (He et al. 2018; Rebuffi et al. 2017; Castro et al. 2018). Another simpler setting, known as the task-incremental setting, involves learning disjoint classifiers for each task.

In this paper, we proposed a framework for the few-shot class-incremental learning (FSCIL) problem. The incremental nature of training makes the few-shot learning (FSL) problem even more challenging. On the other hand, humans can continuously learn new categories from very few samples of data. Therefore, to achieve human-like intelligence, we need to equip deep learning models with the ability to deal with the few-shot class-incremental learning problem.

Training the entire network on classes with very few samples will result in overfitting, which will hamper the network’s performance on test data. Additionally, since the model will not have access to old classes when new classes become available for training, the model will suffer from catastrophic forgetting (French 1999) of the old classes. Therefore, in order to solve the FSCIL problem, we have to address the two issues of overfitting and catastrophic forgetting simultaneously, which makes it even harder.

A common approach of preventing catastrophic forgetting is to ensure that, while training on new classes, the model’s output logits corresponding to the older classes remain unchanged. To achieve this, many methods (Rebuffi et al. 2017; Saihui et al. 2018; Hou et al. 2019; Castro et al. 2018) use a knowledge distillation loss (Hinton, Vinyals, and Dean 2015). The distillation loss can be computed on a few “replay”-ed samples from the old classes. However, the distillation loss is generally biased towards classes with more samples and the new classes. Recently, the authors in (Tao et al. 2020) proposed a method TOPIC to solve the few-shot class-incremental learning problem, using a cognition-based knowledge representation technique. TOPIC uses a neural gas (NG) network (Thomas and Klaus 1991; Fritzke 1995) to model the topology of feature space. While training on new classes, it keeps the topology of NG stable and pushes the samples of new classes to their respective NG node to

\*Equal contribution.

preserve old knowledge.

We propose a novel method, called few-shot lifelong learning (FSLL) for the FSCIL problem by addressing the overfitting and catastrophic forgetting problems from the perspective of the trainable parameters. When a new set of classes becomes available for training, we do not train the entire model on it since the new classes have very few examples, and the full model will quickly overfit to these examples. Instead, we choose very few session trainable parameters to train on these new classes, which reduces the overfitting problem. Our method selects these session trainable parameters in such a way that only unimportant parameters of the model get chosen. As a result, the training on the new set of classes only affects a few unimportant model parameters. Since the important parameters in the model are not affected, the model can retain the old knowledge, thereby minimizing catastrophic forgetting. We encourage the session trainable parameters to be properly updated, but not deviate far from their previous values. To ensure this, we add a regularization loss on the session trainable parameters. Additionally, we maximize the separation between the new and the old class prototypes by minimizing their cosine similarity to improve the network classification performance. We also explore a variant of our method that uses self-supervision as an auxiliary task to improve the model performance further.

We perform experiments on the miniImageNet (Vinyals et al. 2016), CIFAR-100 (Krizhevsky and Hinton 2009), and CUB-200 (Wah et al. 2011) datasets in the FSCIL setting and compare our performance with state-of-the-art method and other baselines. Our experimental results show the effectiveness of our method. We also perform extensive ablation experiments to validate the components of our method. Our main contributions are as follows:

- We propose a novel method for the few-shot class-incremental learning problem. Our proposed method selects very few unimportant model parameters to train every new set of classes in order to minimize overfitting and the catastrophic forgetting problem.
- We empirically show that using self-supervision as an auxiliary task can further improve the performance of the model in this setting.
- We experimentally show that our proposed method significantly outperforms all baselines and the state-of-the-art methods for all the compared datasets.

## Proposed Method

### Problem Setting

In the FSCIL setting, we have a sequence of labeled training sets  $D^{(1)}, D^{(2)}, \dots$ , where  $D^{(t)} = \{(\mathbf{x}_j^{(t)}, y_j^{(t)})\}_{j=1}^{|D^{(t)}|}$ .  $L^{(t)}$  represents the set of classes of the  $t^{\text{th}}$  training set, where  $\forall i, j, L^{(i)} \cap L^{(j)} = \emptyset$  for  $i \neq j$ . The first training set,  $D^{(1)}$ , consist of base classes with a reasonably large number of training examples per class. The remaining training sets  $D^{(t>1)}$  are few-shot training sets of new classes, where each class has very few training samples. The model has to be incrementally trained on  $D^{(1)}, D^{(2)}, \dots$  and only




Figure 1: Our proposed few-shot lifelong learning method. We initially train the network on the base training set  $D^{(1)}$  that contains many examples per class. In the first session, all the parameters are trainable (marked in green). After finishing the training on  $D^{(1)}$ , we find the important (marked in deep blue) and unimportant parameters (marked in light blue) in the model. For training on the few-shot training set  $D^{(t>1)}$ , we select a few unimportant parameters as the session trainable parameters (marked in green). After completing the training on each few-shot training set  $D^{(t>1)}$ , we re-identify the important and unimportant parameters and select the few session trainable parameters for the next session. By preserving the important parameters in the model, the model can preserve the old knowledge. Further, by training only a few session trainable parameters for each few-shot training set, overfitting is also reduced.

$D^{(t)}$  is available at the  $t^{\text{th}}$  training session. After training on  $D^{(t)}$ , the model is evaluated on all the encountered classes in  $L^{(1)}, \dots, L^{(t)}$ . Each of the few-shot training sets  $D^{(t>1)}$ , contain  $C$  classes and  $K$  training examples per class. This setting is referred a  $C$ -way  $K$ -shot FSCIL setting.

Since every few-shot training set has very few examples per class, storing examples from such classes will effectively violate the incremental learning setting. Therefore, our proposed method does not store any examples from the previously seen classes. Also, since FSCIL is based on the class-incremental setting, no task information is available during test time, and the model has to perform classification jointly for all the seen classes.

## Method Overview

In the FSCIL setting, the naive approach is to incrementally train on the training sets  $D^{(1)}, D^{(2)}, \dots$ . However, this approach will lead to a catastrophic forgetting of the older classes. Additionally, since the training set  $D^{(t>1)}$  has very few training examples per class, the model will overfit to these few examples. We propose a novel method to deal with these two challenges.

Our network consists of a feature extractor  $\Theta_F$  and a fully connected classifier  $\Theta_C$ . In our approach, we first train the complete network on the base training set  $D^{(1)}$  for classification using the cross-entropy loss similar to TOPIC. This is a common practice in the few-shot learning setting. During this session, all the parameters of the network are trainable (Fig. 1).

$$L_{D^{(1)}}(\mathbf{x}, y) = F_{CE}(\Theta_C(\Theta_F(\mathbf{x})), y) \quad (1)$$

where  $\mathbf{x}$  and  $y$  refer to an image and its label,  $(\mathbf{x}, y) \in D^{(1)}$ ,  $F_{CE}$  refers to the cross entropy loss.

We discard  $\Theta_C$  after completing the training on  $D^{(1)}$ . Using the trained feature extractor  $\Theta_F$ , we extract the features of all the training samples of  $D^{(1)}$  and average them class-wise to obtain the class prototypes of the base classes.

$$Pr[c] = \frac{1}{N_c} \sum_{k=1}^N \mathbb{I}_{(y_k=c)}(\Theta_F(\mathbf{x}_k)) \quad (2)$$

where  $Pr[c]$  is the prototype of class  $c$  in  $D^{(t)}$ ,  $N_c$  is the number of training examples in the class  $c$ ,  $N$  is the number of training examples in  $D^{(t)}$ ,  $\forall k (\mathbf{x}_k, y_k) \in D^{(t)}$ .  $\mathbb{I}_{(y_k=c)}$  is an indicator function that returns 1 only when the label of the sample  $x_k$  is the class  $c$ , otherwise it returns 0.

The session 1 test examples belong to all the classes encountered till now, i.e., all the base training set classes. For each test example, in session 1, we find the nearest class prototype and predict that class as the output class.

When the session  $t > 1$  starts, the  $D^{(t>1)}$  training set becomes available and data from all previous training sets  $\{D^{(1)}, D^{(2)}, \dots, D^{(t-1)}\}$  become inaccessible. We select very few unimportant parameters from  $\Theta_F$  for training on the few-shot classes in session  $t$  (Fig. 1). We refer to these parameters as the session trainable parameters  $P_{ST}^t$  for session  $t$ . Parameters/weights with low absolute magnitude contribute very less to the final accuracy and are, therefore, unimportant (Han et al. 2015). In order to select the session trainable parameters, we choose a threshold for each layer in  $\Theta_F$ . All parameters in a layer having absolute value lower than the threshold are chosen as the session trainable parameters  $P_{ST}^t$  for session  $t$ . We provide an ablation to show the effect of the proportion of session trainable parameters on the final accuracy. Since we choose the parameters with the lowest absolute weight values, it is highly unlikely that the ‘‘high’’ importance/absolute weight valued parameter will get selected as a less important parameter for the subsequent tasks.

We refer to the remaining parameters as the knowledge retention parameters  $P_{KR}^t$  for session  $t$ , and we keep them

frozen during this session. Since we choose only the unimportant parameters for training, the important parameters remain intact in the model. Therefore, our approach prevents the loss of knowledge from the previously seen classes and reduces catastrophic forgetting.

We train the session trainable parameters  $P_{ST}^t$  on a triplet loss, in order to bring examples from the same class closer and push away those from different classes.

$$L_{TL}(x_i, x_j, x_k) = \max(d(\Theta_F(x_i), \Theta_F(x_j)) - d(\Theta_F(x_i), \Theta_F(x_k)), 0) \quad (3)$$

where  $x_i, x_j, x_k$  are the images in  $D^{(t>1)}$ ,  $L_{TL}$  refers to the triplet loss,  $d$  refers to euclidean distance. Let  $y_i, y_j, y_k$  be the corresponding class labels of  $x_i, x_j, x_k$  and  $y_i = y_j, y_i \neq y_k$ .

We encourage the session trainable parameters to be properly updated, but not deviate far from their previous values. We apply a regularization loss on  $P_{ST}^t$  to achieve this goal. For the regularization loss, we use  $\ell_1$ -regularization between the current  $P_{ST}^t$  parameters weights and their previous values.

$$L_{RL} = \sum_{i=1}^{N_p^t} ||w_i^t - w_i^{t-1}||_1 \quad (4)$$

where  $N_p^t$  refers to the number of session trainable parameters  $P_{ST}^t$  for the training set  $t$ .  $w_i^t, w_i^{t-1}$  refer to the current and previous weights of the  $i^{th}$  parameter in  $P_{ST}^t$ .

Additionally, we apply a cosine similarity loss to minimize the similarity between the prototypes of the older classes  $Pr^{prev}$  and those of the new classes  $Pr^t$ . The new class prototypes are computed using Eq. 2 for  $D^{(t)}$ .

$$L_{CL} = \sum_{i=1}^{N_{Pr}^t} \sum_{j=1}^{N_{Pr}^{prev}} F_{cos}(Pr^t[i], Pr^{prev}[j]) \quad (5)$$

where  $Pr^t$  refers to the prototypes of  $D^{(t)}$ ,  $Pr^{prev}$  refers to set of prototypes from all the previous classes.  $N_{Pr}^t, N_{Pr}^{prev}$  refer to the number of class prototypes in the current training set and all the previous training sets respectively.  $F_{cos}$  refers to the cosine distance loss.  $Pr^t[i], Pr^{prev}[j]$  refer to the  $i^{th}$  and  $j^{th}$  prototypes in  $Pr^t$  and  $Pr^{prev}$ , respectively.

Therefore, the total loss for the training set  $D^{(t>1)}$  is as follows:

$$L(D^{(t>1)}) = L_{TL} + L_{CL} + \lambda L_{RL} \quad (6)$$

where  $\lambda$  is a hyper-parameter that determines the contribution of the regularization loss.

After completing the training on  $D^{(t>1)}$ , we extract the features of the training samples of all the classes in the current training set using the trained feature extractor  $\Theta_F$  and compute the class-wise mean/prototype of these features (Eq. 2). We perform the nearest prototype-based classification using the prototypes of all classes to predict the nearest class for each test example in the current session.

## Self-Supervised Auxiliary Task

We also experiment with a variant of our method, where we train the complete network on  $D^{(1)}$  using the standard cross-entropy loss and an auxiliary self-supervision loss. We use rotation prediction as our auxiliary task (Gidaris, Singh, and Komodakis 2018).

In order to add the auxiliary task, we add a rotation prediction network  $\Theta_R$  after  $\Theta_F$ , in parallel with  $\Theta_C$ . We rotate the each training sample in  $D^{(1)}$  by either 0, 90, 180, and 270 degrees and train the network to predict the angle of rotation using  $\Theta_R$ . The image feature extracted by  $\Theta_F$  is given to  $\Theta_R$  for the rotation prediction task. The total loss for training on  $D^{(1)}$  in this case is as follows:

$$L_{D^{(1)}}(\mathbf{x}, y) = F_{CE}(\Theta_C(\Theta_F(\mathbf{x})), y) + F_{CE}(\Theta_R(\Theta_F(\mathbf{x})), y^r) \quad (7)$$

where  $(\mathbf{x}, y) \in D^{(1)}$  and  $y^r$  is the angle of rotation that  $\mathbf{x}$  was rotated by.

We empirically show that the performance of our method can be improved further using self-supervision. In the ablation studies section, we experimentally show that the rotation prediction-based self-supervision task performs better than SimClr and patch location prediction methods when used as an auxiliary task in our method.

## Related Work

The lifelong/continual learning problem can have two settings: class-incremental setting and task-incremental setting.

### Class-Incremental Lifelong Learning

Class-incremental lifelong learning involves training a model on multiple sets of disjoint classes in a sequence and testing on all the encountered classes. iCaRL, is a popular method proposed in (Rebuffi et al. 2017), which stores class exemplars and learns using the nearest neighbor classification loss on the new classes and a distillation loss on the old class exemplars. The work in (Castro et al. 2018) proposes EEIL, which trains the model using a cross-entropy loss and a distillation loss. NCM (Hou et al. 2019) uses cosine distance metric to reduce the bias of the model towards the new classes. Similarly, BIC (Yue et al. 2019) learns a correction model to reduce the bias in the output logits. We focus on class-incremental learning but in a few-shot setting, which is a more challenging problem due to the few-shot nature of the classes.

### Task-Incremental Lifelong Learning

Task-incremental lifelong learning involves training a model on multiple tasks (which are disjoint sets of classes) in a sequence but maintaining a separate classifier for each task. As a result of a reduced search space, this setting is simpler than the class-incremental setting. Task-incremental lifelong learning methods can be of three types: a) regularization-based, b) replay-based, and c) dynamic network-based.

Regularization-based methods try to reduce changes in the output logits/ important parameters of the network while

training on new tasks in order to preserve the old task knowledge (Lee et al. 2017; Zenke, Poole, and Ganguli 2017; Liu et al. 2018). The work in (Li and Hoiem 2018) uses knowledge distillation to achieve this goal. EWC (Kirkpatrick et al. 2017) decreases the learning rate for the parameters that are important to the older tasks.

Replay-based methods (Lopez-Paz et al. 2017; Chaudhry et al. 2018), store exemplars from old tasks and include them in the training process of the new tasks in order to reduce catastrophic forgetting. Some methods utilize generative models to generate data for the old tasks instead of storing the exemplars (Shin et al. 2017; Wu et al. 2018; Zhai et al. 2019; Xiang et al. 2019).

Dynamic network-based methods modify the network to train on new tasks (Mallya and Lazebnik 2018; Mallya, Davis, and Lazebnik 2018; Aljundi et al. 2018; Serrà et al. 2018; Yoon et al. 2017). These methods employ techniques such as dynamic expansion, network pruning, and parameter masking to prevent catastrophic forgetting. PackNet, a method proposed in (Mallya and Lazebnik 2018), utilizes pruning to free parameters for training new tasks. The work in (Serrà et al. 2018) proposes to learn attention masks for old tasks to constrain the parameters when training on the new task. The authors in (Xu and Zhu 2018) utilize reinforcement learning to decide the number of additional neurons needed for each new task.

Since we focus on the class incremental setting in this paper, the task-incremental methods do not apply to this setting. Therefore, we have to exclude them for comparison in the experimental section.

### Few-Shot Learning

Few-shot learning (FSL) methods train models to perform well for classes with very few training examples (few-shot classes). Many research works deal with the few-shot learning problem. Few-shot learning methods generally employ meta-learning and metric learning techniques (Vinyals et al. 2016; Snell, Swersky, and Zemel 2017; Sung et al. 2018; Finn, Abbeel, and Levine 2017; Sun et al. 2019). However, most of them are focused solely on the few-shot classes. Recently some methods have also explored the loss of performance in the non few-shot classes due to the techniques used to benefit the few-shot classes (Gidaris and Komodakis 2018; Ren et al. 2019). The method proposed in (Gidaris and Komodakis 2018) extends an object recognition system with an attention-based few-shot classification weight generator and redesigns the classifier as a similarity function between feature representations and classification weight vectors. It combines the recognition of both the few-shot and non few-shot classes.

Most of the standard few-shot learning methods perform testing on few-shot episodes, containing a few classes with very few labeled samples. By reducing the search space to the few classes present in the episode, the problem becomes much simpler. On the other hand, the few-shot class-incremental learning setting performs testing on all the encountered classes, which is more realistic and challenging. TOPIC, proposed in (Tao et al. 2020), utilizes a neural gas network (Thomas and Klaus 1991; Fritzke 1995) to

model the topology of feature space and stabilizes the topology while introducing new classes in order to preserve old knowledge. This method achieves state-of-the-art results in the FSCIL setting, and we have compared our results with this method.

## Self-Supervised Learning

While obtaining labeled data is expensive and time-consuming, recent work has considered alternative mechanisms that can substitute for such explicitly labeled supervision. In particular, the self-supervised learning paradigm trains a network on data using labels extracted from the data itself. Self-supervised learning helps the network to learn better features. To perform self-supervised learning, various types of pseudo tasks/labels are used to train the network, such as image inpainting or image completion (Pathak et al. 2016), image colorization (Larsson, Maire, and Shakhnarovich 2016; Zhang, Isola, and Efros 2016), prediction of relative patch position (Doersch, Gupta, and Efros 2015), solving Jigsaw puzzles in (Noroozi and Favaro 2016). In (Gidaris, Singh, and Komodakis 2018), the authors propose to rotate the images by a fixed set of angles, and the network is trained to predict the angle of rotation. It is a very popular method for self-supervision.

Contrastive Multiview Coding (CMC) (Tian, Krishnan, and Isola 2019) trains the network to maximize the mutual information between the different views of an image but requires a specialized architecture, including separate encoders for different views of the data. Momentum Contrast (MoCo) (He et al. 2020) matches encoded queries  $q$  to a dictionary of encoded keys using a contrastive loss, but it requires a memory bank to store the dictionary. SimCLR (Chen et al. 2020) augments the input to produce two different but correlated views and uses contrastive loss to bring them closer in the feature space. It does not require specialized architectures or a memory bank and still achieves state-of-the-art unsupervised learning results, outperforming the CMC and MoCo self-supervision techniques.

## Experiment

In this section, we describe the datasets and implementation details of the experiments that we conduct.

### Datasets

We perform experiments in the FSCIL setting using three image classification datasets CIFAR-100 (Krizhevsky and Hinton 2009), miniImageNet (Vinyals et al. 2016) and CUB-200 (Wah et al. 2011). The CIFAR-100 dataset consists of 100 classes with each class containing 500 training images and 100 testing images. Each of the 60,000 images is of size  $32 \times 32$ . The miniImageNet dataset also consists of 60,000 images from 100 classes, chosen from the ImageNet-1k dataset (Deng et al. 2009). There are 500 training and 100 test images of size  $84 \times 84$  for each class. The CUB-200 dataset consists of about 6,000 training images and 6,000 test images for 200 categories of birds. The images are resized to  $256 \times 256$  and then cropped to  $224 \times 224$  for training.

In the case of the CIFAR-100 and miniImageNet datasets, we choose 60 and 40 classes as the base and new classes, respectively. For every few-shot training set, we use a 5-way 5-shot setting, i.e., each few-shot training set has 5 classes with 5 training examples per class. Therefore, we have 1 base training set and 8 few-shot training sets (total 9 training sessions) for the CIFAR-100 and miniImageNet datasets. For the CUB-200 dataset, we choose 100 and 100 classes as the base and new classes, respectively. For every few-shot training set of CUB-200, we use a 10-way 5-shot setting, i.e., each few-shot training set has 10 classes with 5 training examples per class. Therefore, we have 1 base training set and 10 few-shot training sets (total 11 training sessions) for the CUB-200 dataset. We construct each few-shot training set by randomly choosing 5 training examples per class, while the test set contains test examples from all the encountered classes. For a fair comparison, we used the same dataset settings as used in (Tao et al. 2020).

### Implementation Details

We use ResNet-18 (He et al. 2015) architecture for our experiments on all the three datasets. The last classification layer of ResNet-18 is  $\Theta_C$ , and the remaining network serves as the feature extraction network  $\Theta_F$ . We train  $\Theta_F^{(1)}$ , and  $\Theta_C^{(1)}$  on the base training set  $D^{(1)}$  with an initial learning rate of 0.1 and mini-batch size of 128. After the 30 and 40 epochs, we reduce the learning rate to 0.01 and 0.001, respectively. We train on  $D^{(1)}$  for a total of 50 epochs and then discard  $\Theta_C^{(1)}$ .

We finetune the feature extractor on each of the few-shot training sets  $D^{(t>1)}$  for 30 epochs, with a learning rate of  $1e-4$  (and  $1e-3$  for CUB-200). We set the threshold values for each layer in such a way that only 10% of  $\Theta_F$  get selected as the session trainable parameters in all our experiments. Since the few-shot training sets contain very few training examples, the mini-batch contains all the examples. After training the feature extractor on  $D^{(t)}$ , we test  $\Theta_F^t$  on the combined test sets of all encountered classes. Session  $t$  accuracy refers to the total accuracy over all the the classes encountered till that session ( $L^{(1)}, L^{(2)} \dots, L^{(t)}$ ).

We perform standard random cropping and flipping for data augmentation proposed in (He et al. 2015; Hou et al. 2019) for all methods. Since we have very few new class training samples, we use the batchnorm statistics computed on  $D^{(1)}$  and fix the batchnorm layers while finetuning on  $D^{(t>1)}$  as done in (Tao et al. 2020). We run each experiment 10 times and report the average test accuracy over all the encountered classes. The standard deviations among the runs are low (around 0.5% on average) for all the experiments. For the experiments with the auxiliary self-supervision task, we use a convolutional neural network as  $\Theta_R$  to predict the rotation angle.  $\Theta_R$  consists of 4 convolutional layers, each containing 512 filters of filter size of 3, stride 1, and padding 1. We use an adaptive average pooling layer and a linear layer of output size 4 after the last convolutional layer. For reporting the results, we also include the base set  $D^{(1)}$  accuracy for a fair comparison with TOPIC and other methods.

Method	Sessions											Our Relative Improvements
	1	2	3	4	5	6	7	8	9	10	11	
Ft-CNN (Tao et al. 2020)	68.68	44.81	32.26	25.83	25.62	25.22	20.84	16.77	18.82	18.25	17.18	<b>+28.37</b>
Joint-CNN (Tao et al. 2020)	68.68	62.43	57.23	52.80	49.50	46.10	42.80	40.10	38.70	37.10	35.60	<b>+9.95</b>
iCaRL (Rebuffi et al. 2017)	68.68	52.65	48.61	44.16	36.62	29.52	27.83	26.26	24.01	23.89	21.16	<b>+24.39</b>
EEIL (Castro et al. 2018)	68.68	53.63	47.91	44.20	36.30	27.46	25.93	24.70	23.95	24.13	22.11	<b>+23.44</b>
NCM (Hou et al. 2019)	68.68	57.12	44.21	28.78	26.71	25.66	24.62	21.52	20.12	20.06	19.87	<b>+25.68</b>
TOPIC (Tao et al. 2020)	68.68	62.49	54.81	49.99	45.25	41.40	38.35	35.36	32.22	28.31	26.28	<b>+19.27</b>
<b>FSSL (Ours)</b>	<b>68.72</b>	<b>65.67</b>	<b>62.33</b>	<b>58.10</b>	<b>55.44</b>	<b>52.66</b>	<b>51.17</b>	<b>50.27</b>	<b>48.31</b>	<b>47.25</b>	<b>45.55</b>	<b>0</b>
<b>FSSL* (Ours)</b>	<b>72.77</b>	<b>69.33</b>	<b>65.51</b>	<b>62.66</b>	<b>61.10</b>	<b>58.65</b>	<b>57.78</b>	<b>57.26</b>	<b>55.59</b>	<b>55.39</b>	<b>54.21</b>	-
<b>FSSL*+SS (Ours)</b>	<b>75.63</b>	<b>71.81</b>	<b>68.16</b>	<b>64.32</b>	<b>62.61</b>	<b>60.10</b>	<b>58.82</b>	<b>58.70</b>	<b>56.45</b>	<b>56.41</b>	<b>55.82</b>	-

Table 1: Results on CUB-200 using the ResNet-18 architecture on the 10-way 5-shot FSCIL setting. We compare our method with TOPIC (CVPR’20) which is the state-of-the-art method for this setting. Session  $t$  accuracy refers to the total accuracy over all the the classes encountered till that session ( $L^{(1)}, L^{(2)} \dots, L^{(t)}$ ).




Figure 2: Results on miniImageNet using the ResNet-18 architecture on the 5-way 5-shot FSCIL setting

### Baselines and Compared Methods

We compare our method with iCaRL (Rebuffi et al. 2017), EEIL (Castro et al. 2018) and NCM (Hou et al. 2019) in the FSCIL setting as in (Tao et al. 2020). We compare our method with the Ft-CNN, which involves only finetuning the model on the few training examples of  $D^{(t>1)}$ . We also compare our method with the Joint-CNN method, which trains on the combined data of the base and few-shot classes.

### CUB-200 Results

The results in Table 1 indicate that our method significantly outperforms the Ft-CNN model on CUB-200. Our method performs significantly better than the Joint-CNN. This is because CUB-200 contains 100 few-shot classes in this setting and the Joint-CNN model overfits to these classes, resulting in lower overall performance. Our method outperforms the state-of-the-art TOPIC method by an absolute margin of 19.27%. Even if we exclude the base training set ( $D^{(1)}$ ) accuracy, our method achieves an average accuracy of 27% on the few-shot training sets.

While training on the model on  $D^{(1)}$ , we observed that using an initial learning rate of 0.01 achieves a better session 1 accuracy than reported in (Tao et al. 2020). For complete-




Figure 3: Results on CIFAR-100 using the ResNet-18 architecture on the 5-way 5-shot FSCIL setting.

ness, we also provide the results for this model (FSSL\*). We perform an additional experiment (FSSL\*+SS), where we also train the network on an auxiliary self-supervised rotation prediction task during the training of  $D^{(1)}$ .

### miniImageNet Results

Fig. 2 depicts the performance of different methods on the miniImageNet FSCIL setting. Our method significantly outperforms the Ft-CNN model and performs slightly better than the Joint-CNN model because the Joint-CNN model overfits due to the presence of many few-shot classes. Our method significantly outperforms the state-of-the-art TOPIC model by around 15.07%. We observe that the performance can be improved further by tuning the weight decay option of the SGD optimizer, which we take as  $1e-3$  (FSSL\*).

### CIFAR-100 Results

Fig. 3 depicts the performance of different methods on the CIFAR-100 FSCIL setting. Our method significantly outperforms the state-of-the-art TOPIC model, by an absolute margin of 9.09%. We perform an additional experiment (FSSL+SS), where we have also trained the network on an auxiliary self-supervised rotation prediction task during the training of  $D^{(1)}$ .

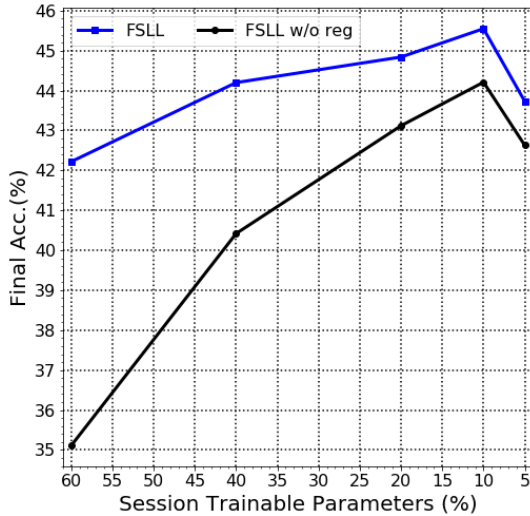


Figure 4: Performance of FSSL on the CUB-200 FSCIL setting with and without regularization using different proportions of session trainable parameters.

## Ablation Experiments

We perform various ablation experiments to validate our method.

### Proportion of Session Trainable Parameters

Fig. 4 shows the effect of increasing the proportion of session trainable parameters on the final accuracy for CUB-200. The final accuracy in Fig. 4 refers to the session 11 test result ( $S_{11}$ ). If we choose a high threshold for selecting unimportant parameters, then there will be a high proportion of session trainable parameters, and it may include high absolute value/important parameters. We observe that FSSL performance suffers in such a case. FSSL achieves the best performance when the session trainable parameters are 10% of  $\Theta_F$ . As we decrease the proportion of session trainable parameters, the proportion of knowledge retention parameters increases, and the performance of the model improves till we reach 10%. If we choose less than 10% of  $\Theta_F$  as the session trainable parameters, the model performance starts dropping due to the shortage of trainable parameters (underfitting).

### Significance of Regularization

Fig. 4 shows the effect of removing the regularization loss from our method. When the proportion of session trainable parameters is high, the corresponding proportion of knowledge retaining parameters is low, and therefore, the regularization loss plays a critical role in the model performance. Even when the proportion of session trainable parameters is low ( $\approx 10\%$ ), the regularization loss provides improvement to the performance, as shown in Fig. 4.

### Choice of Regularization Hyper-Parameter

Table 2, reports the effect of changing the regularization hyper-parameter  $\lambda$  on the performance of the model for the

$\lambda$	1	3	5	7	9
$S_{11}$	44.66%	44.83%	<b>45.55%</b>	44.96%	44.87%

Table 2: Session 11 ( $S_{11}$ ) classification results on CUB-200 using the ResNet-18 architecture on the 10-way 5-shot FSCIL setting for different values for the regularization hyper-parameter  $\lambda$  values.

Auxiliary SS	Patch	SimClr	Rotation	w/o SS
$S_{11}$	54.56%	54.71%	55.82%	54.21%

Table 3: Session 11 classification results on CUB-200 using FSSL\* for the 10-way 5-shot FSCIL setting for different types of auxiliary self-supervised (SS) task.

CUB-200 dataset in the FSCIL setting. We have reported the session 11 ( $S_{11}$ ) test results in this table. We observe the best model performance for  $\lambda = 5$ , and we use this value of the regularization hyper-parameter for all our experiments.

### Significance of Cosine Similarity Loss

We performed ablations to verify the significance of the prototype cosine similarity loss. We observe that in the absence of the prototype cosine similarity loss, the session 11 model performance ( $S_{11}$ ) for the CUB-200 dataset drops from 45.55% to 44.32%.

### Choice of Self-Supervised Auxiliary Task

We perform experiments with the auxiliary task as relative patch location prediction (Doersch, Gupta, and Efros 2015) (Patch), rotation angle prediction (Rotation) (Gidaris et al. 2019) and SimCLR (Chen et al. 2020). SimCLR utilizes contrastive learning and is the state-of-the-art self-supervision technique. Table 3, shows that the rotation-based auxiliary self-supervised task performs significantly better than patch and SimCLR methods.

### Self-Supervision for Few-Shot Classes

We also perform experiments to include the self-supervised auxiliary task in the training process for the few-shot training sets ( $D^{(t>1)}$ ) along with the base training set  $D^{(1)}$ . Our experiments on the CUB-200 dataset show that this results in a session 11 ( $S_{11}$ ) accuracy of 54.34%, which is lower than 55.82% achieved by FSSL\*+SS. Therefore, using the self-supervised auxiliary task for training on  $D^{(t>1)}$  does not produce any benefits.

## Conclusion

We propose a novel Few-Shot Lifelong Learning (FSSL) method for the few-shot class-incremental learning problem. Our method selects very few unimportant parameters as the session trainable parameters to train on every new set of few-shot classes to deal with the problems of overfitting and catastrophic forgetting. We empirically show that FSSL significantly outperforms the state-of-the-art method. We experimentally show that using self-supervision as an auxiliary task can further improve the performance of the model in this setting.

## References

- Aljundi, R.; Babiloni, F.; Elhoseiny, M.; Rohrbach, M.; and Tuytelaars, T. 2018. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 139–154.
- Castro, F. M.; Marín-Jiménez, M. J.; Guil, N.; Schmid, C.; and Alahari, K. 2018. End-to-end incremental learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 233–248.
- Chaudhry, A.; Ranzato, M.; Rohrbach, M.; and Elhoseiny, M. 2018. Efficient lifelong learning with a-gem. *arXiv preprint arXiv:1812.00420*.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv preprint arXiv:2002.05709*.
- Deng, J.; Dong, W.; Socher, R.; Li, L. J.; Li, K.; and Li, F. F. 2009. ImageNet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, 248–255.
- Doersch, C.; Gupta, A.; and Efros, A. A. 2015. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision*, 1422–1430.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 1126–1135. JMLR. org.
- French, R. M. 1999. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences* 3(4): 128–135.
- Fritzke, B. 1995. A Growing Neural Gas Network Learns Topologies. *Advances in Neural Information Processing Systems* 7.
- Gidaris, S.; Bursuc, A.; Komodakis, N.; Pérez, P.; and Cord, M. 2019. Boosting few-shot visual learning with self-supervision. In *Proceedings of the IEEE International Conference on Computer Vision*, 8059–8068.
- Gidaris, S.; and Komodakis, N. 2018. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4367–4375.
- Gidaris, S.; Singh, P.; and Komodakis, N. 2018. Unsupervised Representation Learning by Predicting Image Rotations. In *International Conference on Learning Representations*.
- Han, S.; Pool, J.; Tran, J.; and Dally, W. 2015. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, 1135–1143.
- He, C.; Wang, R.; Shan, S.; and Chen, X. 2018. Exemplar-Supported Generative Reproduction for Class Incremental Learning. In *Proceedings of the British Machine Vision Conference*.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *Computer Science* 14(7): 38–39.
- Hou, S.; Pan, X.; Loy, C. C.; Wang, Z.; and Lin, D. 2019. Learning a Unified Classifier Incrementally via Rebalancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 831–839.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114(13): 3521–3526.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. Technical report, Citeseer.
- Larsson, G.; Maire, M.; and Shakhnarovich, G. 2016. Learning representations for automatic colorization. In *European Conference on Computer Vision*, 577–593. Springer.
- Lee, S.-W.; Kim, J.-H.; Jun, J.; Ha, J.-W.; and Zhang, B.-T. 2017. Overcoming catastrophic forgetting by incremental moment matching. In *Advances in Neural Information Processing Systems*, 4652–4662.
- Li, Z.; and Hoiem, D. 2018. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40(12): 2935–2947.
- Liu, X.; Masana, M.; Herranz, L.; Joost, V. D. W.; Lopez, A. M.; and Bagdanov, A. D. 2018. Rotate your Networks: Better Weight Consolidation and Less Catastrophic Forgetting. *arxiv preprint arXiv:1802.02950*.
- Lopez-Paz, D.; et al. 2017. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 6467–6476.
- Mallya, A.; Davis, D.; and Lazebnik, S. 2018. Piggyback: Adapting a single network to multiple tasks by learning to mask weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 67–82.
- Mallya, A.; and Lazebnik, S. 2018. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7765–7773.
- Noroozi, M.; and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision*, 69–84. Springer.
- Pathak, D.; Krahenbuhl, P.; Donahue, J.; Darrell, T.; and Efros, A. A. 2016. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2536–2544.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2001–2010.



- Ren, M.; Liao, R.; Fetaya, E.; and Zemel, R. 2019. Incremental few-shot learning with attention attractor networks. In *Advances in Neural Information Processing Systems*, 5276–5286.
- Saihui, H.; Xinyu, P.; Chen Change, L.; Zilei, W.; and Dahua, L. 2018. Lifelong learning via progressive distillation and retrospection. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Serrà, J.; Suris, D.; Miron, M.; and Karatzoglou, A. 2018. Overcoming catastrophic forgetting with hard attention to the task. *arXiv preprint arXiv:1801.01423*.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. In *Advances in Neural Information Processing Systems*, 2990–2999.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, 4077–4087.
- Sun, Q.; Liu, Y.; Chua, T.-S.; and Schiele, B. 2019. Meta-transfer learning for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 403–412.
- Sung, F.; Yang, Y.; Zhang, L.; Xiang, T.; Torr, P. H.; and Hospedales, T. M. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1199–1208.
- Tao, X.; Hong, X.; Chang, X.; Dong, S.; Wei, X.; and Gong, Y. 2020. Few-Shot Class-Incremental Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12183–12192.
- Thomas, M.; and Klaus, S. 1991. A "Neural-Gas" Network Learns Topologies. *Artificial Neural Networks*.
- Tian, Y.; Krishnan, D.; and Isola, P. 2019. Contrastive Multiview Coding. *arXiv preprint arXiv:1906.05849*.
- Vinyals, O.; Blundell, C.; Lillicrap, T.; Kavukcuoglu, K.; and Wierstra, D. 2016. Matching Networks for One Shot Learning. *arXiv preprint arXiv:1606.04080*.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology.
- Wu, C.; Herranz, L.; Liu, X.; van de Weijer, J.; Raducanu, B.; et al. 2018. Memory replay gans: Learning to generate new categories without forgetting. In *Advances In Neural Information Processing Systems*, 5962–5972.
- Xiang, Y.; Fu, Y.; Ji, P.; and Huang, H. 2019. Incremental Learning Using Conditional Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 6619–6628.
- Xu, J.; and Zhu, Z. 2018. Reinforced continual learning. In *Advances in Neural Information Processing Systems*, 899–908.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2017. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*.
- Yue, W.; Yinpeng, C.; Lijuan, W.; Yuancheng, Y.; Zicheng, L.; Yandong, G.; and Yun, F. 2019. Large Scale Incremental Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 3987–3995. JMLR. org.
- Zhai, M.; Chen, L.; Tung, F.; He, J.; Nawhal, M.; and Mori, G. 2019. Lifelong gan: Continual learning for conditional image generation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2759–2768.
- Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *European Conference on Computer Vision*, 649–666. Springer.