

Covariance-based Space Regularization for Few-shot Class Incremental Learning

Yijie Hu¹, Guanyu Yang¹, Zhaorui Tan¹, Xiaowei Huang², Kaizhu Huang³, Qiu-Feng Wang¹
 Xi'an Jiaotong-Liverpool University¹, Liverpool University², Duke Kunshan University³

{Yijie.Hu20, Zhaorui.Tan21}@student.xjtlu.edu.cn, {Guanyu.Yang02, Qiufeng.Wang}@xjtlu.edu.cn
 xiaowei.huang@liverpool.ac.uk, kaizhu.huang@dukekunshan.edu.cn

Abstract

Few-shot Class Incremental Learning (FSCIL) presents a challenging yet realistic scenario, which requires the model to continually learn new classes with limited labeled data (i.e., incremental sessions) while retaining knowledge of previously learned base classes (i.e., base sessions). Due to the limited data in incremental sessions, models are prone to overfitting new classes and suffering catastrophic forgetting of base classes. To tackle these issues, recent advancements resort to prototype-based approaches to constrain the base class distribution and learn discriminative representations of new classes. Despite the progress, the limited data issue still induces ill-divided feature space, leading the model to confuse the new class with old classes or fail to facilitate good separation among new classes. In this paper, we aim to mitigate these issues by directly constraining the span of each class distribution from a covariance perspective. In detail, we propose a simple yet effective covariance constraint loss to force the model to learn each class distribution with the same covariance matrix. In addition, we propose a perturbation approach to perturb the few-shot training samples in the feature space, which encourages the samples to be away from the weighted distribution of other classes. Regarding perturbed samples as new class data, the classifier is forced to establish explicit boundaries between each new class and the existing ones. Our approach is easy to integrate into existing FSCIL approaches to boost performance. Experiments on three benchmarks validate the effectiveness of our approach, achieving state-of-the-art performance.

1. Introduction

Provided with a substantial amount of stationary data, recent advancements in deep learning have enabled neural networks to excel in classification tasks [9, 17, 30, 31, 35]. However, it is not feasible to directly deploy neural net-

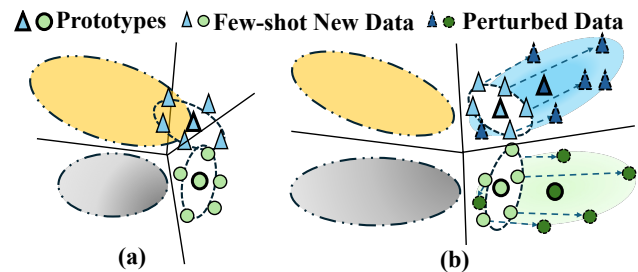


Figure 1. (a) Prototype-based models demonstrate compact representations of old classes, conserving space for new classes, though risking confusion due to mixed class distributions. (b) The proposed approach aims to regularize each class distribution within a fixed span by constraining covariance and to enhance class separation through the learning of perturbed new class data.

works into the open world, where the data may emerge in a non-stationary way, such as recognizing new types of diseases or new vehicles in autonomous driving. Simply re-training or fine-tuning the model with new data introduces the well-known catastrophic forgetting problem [7, 24]. To address this challenge, Class Incremental Learning (CIL) has been extensively researched to broaden the application scope of neural networks [33, 52]. By simulating the scenario where disjoint new data appears in incremental sessions, CIL aims to learn new concepts without forgetting old knowledge.

The conventional CIL setting assumes the amount of new data is usually sufficient, which may not be realistic as labeling new data can be expensive. To handle this issue, few-shot class incremental learning (FSCIL) has attracted much attention recently, where only a few training samples from new classes are available during incremental sessions [18, 32, 47]. The framework of FSCIL typically involves two key stages. Initially, the model is trained on a base dataset, where all classes (referred to as base classes or old classes) contain sufficient instances. Subsequently, the model engages in incremental sessions, where it is required

to learn new classes with limited samples in each session without access to previous data. After training, the model is evaluated on all previously encountered classes. Due to the scarcity of new data, the model is more vulnerable to overfit new classes and hence suffers catastrophic forgetting of old classes during incremental learning sessions [20, 23, 26, 52].

To overcome both overfitting and catastrophic forgetting issues of FSCIL, recent wisdom resorts to prototype-based models [10, 18, 44, 54]. Specifically, prototype-based models [28, 42] replace the linear classifier with learnable prototypes, aiming to learn the most representative point (e.g., center point) of each class. Recent FSCIL works involve two stages, i.e., base session learning and incremental session learning. In the base learning stage, prototype-based models are utilized to learn compact representations of base classes [10, 25, 44]. During the incremental sessions, the feature extractor is frozen, and new prototypes are computed using the mean of each new class. The frozen feature extractor can effectively alleviate the catastrophic forgetting problem, and the prototype classifier can mitigate the overfitting problem. However, as shown in Fig. 1(a), in the fixed feature space, the estimated new class prototypes may lie very close to base class distributions due to the scarcity of the new data. Thus, class distributions can be ill-located during incremental sessions, i.e., new class distribution may lie close to old classes or mixed with other new class distributions [38, 44]. This dilemma may then lead the model to confuse the class distributions in the feature space. Besides, the computed class prototypes can be affected by data noise, exacerbating the confusion dilemma.

To resolve this dilemma, we argue that each class should occupy the same amount of feature space. Motivated by this, we dynamically relocate each new class during incremental sessions to prevent overlapping. Without loss of generality, we assume each class follows the Gaussian distribution, where the mean controls the location and the covariance controls the scope of each feature distribution. To ensure each class takes up the same amount of feature space, we first regularize the covariance of each class in the base session. However, the direct optimization of the distribution-related statistics during training is often inefficient, as it requires model inferences throughout the entire dataset. Inspired by previous works in variational inference [15, 16, 48], we adopt a similar approach to estimate and regularize the class distributions efficiently. Specifically, we derive an evidence lower bound for classification and distribution learning. Maximizing the lower bound equals to maximize the confidence of prediction and minimize the KL divergence between the estimated distribution and the prior distribution. We set the fixed covariance of the prior distribution to be equal to all classes, which formulates the KL divergence to a covariance constraint loss, encouraging each class distribution to have the same covariance

during training. The covariance constraint loss serves as an upper bound of KL divergence, which leads to a tighter evidence lower bound.

Furthermore, to reallocate the feature space for the few-shot new classes during incremental sessions, we propose a perturbation approach to expand the distributions of few-shot new classes by generating perturbed samples for new classes, then pushing these perturbed samples away from semantically similar classes. In detail, we introduce a learnable prior distribution for each few-shot sample based on semantic similarity. We first obtain each training sample's softmax score towards other classes, then we use the weighted mean and the fixed covariance as a prior distribution. We adopt one linear layer to estimate the mean and variance of the new class distribution. The estimated distribution is supervised by minimizing the KL divergence between the estimated distribution and the prior. Next, we multiply the predicted variance with the extracted features and add the predicted mean to create perturbed samples, as shown in Fig. 1(b). The perturbed samples are treated as new training samples, learning which pushes new classes away from the overlapping distributions. In this manner, the new class distribution is expanded by assigning the fixed variance. Learning semantic-guided perturbed samples also facilitates better separation between classes.

Our approach is easy to implement and can be integrated with other approaches. We choose several recent state-of-the-art FSCIL models as baselines [29, 49] and apply the proposed approach to these methods. Extensive experiments on three benchmarks validate the effectiveness of our approach. Our main contributions can be summarized as follows: (1) We propose a covariance constraint loss (CCL) to regularize the class distributions, constraining the span of each distribution within the same range. (2) We propose a semantic guided perturbation approach to perturb the few-shot new data, aiming to learn extensive and discriminative new class distributions. (3) Our proposed method is easy to incorporate into current FSCIL models to boost their performance. Experimental results on the FSCIL benchmark datasets validate the effectiveness of our approach.

2. Related Works

2.1. Class Incremental Learning

Class Incremental Learning (CIL) aims to learn new classes from a sequence of classification tasks without access to previously encountered data. The primary objective of CIL is to effectively learn new classes while minimizing the forgetting of old classes. Recent research in CIL can be broadly categorized into three approaches. The first and most straightforward method is to retain old data or knowledge during the learning process. Recent works [12, 26, 27, 50, 52] propose to mitigate the forget-

ting issue by rehearsing and generating previously retrained class data. Another common approach involves identifying key model parameters associated with previously learned classes and dynamically updating only the remaining parameters during incremental sessions [14, 20, 41, 45]. The third category focuses on addressing the bias inherent in CIL methods, which tend to favor the most recently learned classes [2, 11, 40].

2.2. Few-Shot Learning

Few-Shot Learning (FSL) aims to develop a classification model with very limited data. To generalize on few-shot classes, metric-based methods focus on learning a similarity metric that can effectively distinguish between classes with minimal examples [28, 34, 37, 46]. Hallucination-based approaches utilize data augmentation techniques, such as geometric transformations, style transfer and statistical augmentations, to increase the amount of training data [3, 8, 39, 43].

2.3. Few-Shot Class Incremental Learning

Few-Shot Class Incremental Learning (FSCIL) aims to learn new classes with limited incoming data in an incremental manner [32, 44, 47, 49, 51]. TOPIC [32] first introduces this setting and employs a neural gas algorithm to preserve the topology in the embedding space. To address the challenge of limited data in incremental sessions, prototype learning [19, 28, 42] has been widely adopted in FSCIL to enhance the model’s generalization to new or unseen data. To mitigate the catastrophic forgetting of old classes, recent studies [10, 44, 47] propose freezing the feature extraction backbone after training the base sessions and computing new class prototypes during incremental sessions. To learn more representative prototypes, Zhu et al. [53] introduces a self-promoted prototype refinement mechanism to develop extensible feature representations in the base session. LDC [21] utilizes a recurrent calibration module to learn new prototypes from sampled data, though this can be inefficient due to its recurrent nature. Unlike previous methods, our approach boosts FSCIL by regularizing the feature space for the few-shot new classes.

3. Methodology

In this section, we first give the problem formulation of FSCIL in Sec. 3.1. We then describe our proposed FSCIL method in detail via two sections, i.e., Covariance Constraint Loss (CCL) in Sec. 3.3 and Semantic Perturbation Learning (SPL) in Sec. 3.4, respectively.

3.1. Preliminaries

3.1.1 Problem Formulation

FSCIL aims to train a classification model with T sequential sessions $\{\mathcal{D}^0, \mathcal{D}^1, \dots, \mathcal{D}^T\}$, where $\mathcal{D}^t = \{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^{|\mathcal{D}^t|}$ is the training dataset at the t -th session. \mathbf{x}_i^t is the i -th input sample and its label $y_i^t \in \mathcal{C}^t$. The label space \mathcal{C}^t of dataset \mathcal{D}^t is disjoint between different sessions, i.e., $\forall t_1 \neq t_2, \mathcal{C}^{t_1} \cap \mathcal{C}^{t_2} = \emptyset$. The first session \mathcal{D}^0 is called the base session, which usually contains a sufficient amount of training samples for each old class $c \in \mathcal{C}^0$. In the next incremental session \mathcal{D}^t , there are N new classes with K training samples (usually 1 or 5 samples) in each class, formulating a N way K shot problem, i.e. $|\mathcal{D}^t| = N \cdot K$. In the t -th session, previous datasets $\{\mathcal{D}^0, \mathcal{D}^1, \dots, \mathcal{D}^{t-1}\}$ are not available, the model can only access the data in \mathcal{D}^t . After training in session t , the model is evaluated on all seen classes $\tilde{\mathcal{C}}^t = \mathcal{C}^0 \cup \mathcal{C}^1 \dots \cup \mathcal{C}^t$.

3.1.2 Prototype-based Model

Researchers [28, 29, 49, 54] commonly adopt the prototype-based framework in FSCIL, where classifier weights are treated as prototypes. During the base session, given the feature extractor $f_\psi : \mathbf{x} \rightarrow x \in \mathbb{R}^d$ and the classifier $\mathbf{W}^0 = \{\mathbf{w}_1^0, \mathbf{w}_2^0 \dots \mathbf{w}_{|\mathcal{C}^0|}^0\}$, the model is trained using the cross-entropy loss:

$$\mathcal{L}_{ce}(\mathbf{x}_i^0, y_i^0) = \mathbb{E}_{(\mathbf{x}_i^0, y_i^0) \sim \mathcal{D}^0} \frac{1}{N} \sum_{c=1}^{|\mathcal{C}^0|} -y_i^0 \log p(y = c | x_i^0), \quad (1)$$

where $x_i^0 = f_\psi(\mathbf{x}_i^0)$ denotes the extracted feature, $p(y = c | x_i^0)$ computes the confidence score of each sample using softmax of the cosine similarities:

$$p(y = c | x_i^0) = \frac{\exp(\cos(\mathbf{w}_c, x_i^0))}{\sum_{c=1}^{|\mathcal{C}^0|} \exp(\cos(\mathbf{w}_c, x_i^0))}. \quad (2)$$

For each incremental session $t > 0$, the feature extractor f_ψ is frozen, and the classifier is updated by adding new class prototypes: $\mathbf{W}^t = \{\mathbf{w}_1^0, \mathbf{w}_2^0 \dots \mathbf{w}_{|\mathcal{C}^0|}^0\} \cup \{\mathbf{w}_1^t \dots \mathbf{w}_{\mathcal{C}^t}^t\}$, where each new class prototype is computed by averaging samples from its corresponding class:

$$\mathbf{w}_c^t = \frac{1}{|\mathcal{D}^t|} \sum_{i=1}^{|\mathcal{D}^t|} f_\psi(\mathbf{x}_i^t). \quad (3)$$

3.2. Overview of the Framework

Our framework follows the two-stage learning procedure in most of recent works [10, 18, 44], which begins with pre-training the model using the base classes. The pretraining stage is called the base session, where we integrate

the proposed covariance constraint loss (CCL) to constrain the learned representations of base classes, as shown in Fig. 2(a). During the incremental sessions, we propose a semantic perturbation learning (SPL) approach, aiming to enlarge the separation between old and new classes by learning the perturbed new data, which is shown in Fig. 2(b).

3.3. Covariance Constraint Loss

Recent wisdom [5, 29, 55] has demonstrated a good pre-trained model benefits the learning of incoming few-shot new classes. The common attempt in recent works is to learn very compact base class representations by pushing data to the learned prototypes via fantasizing new classes or metric-based classification losses [18, 29, 49]. However, simply pushing the data close to the class centroid does not explicitly constrain the span of each class distribution, which may lead to representation confusion when learning new classes (Fig. 1(a)). Hence, we apply a covariance constraint to each base class during training, ensuring distinct means but identical covariance across classes, as shown in Fig. 2(a).

In order to constrain the covariance of each distribution, it is vital to estimate base class distributions $p(x, y)$. However, it is generally computational intractable. To solve this issue, previous works [6, 15, 16] adopt the variational inference, which involves a parametric posterior function $q_\phi(z|x, y)$ to approximate the true posterior by maximizing the evidence lower bound (ELBO):

$$\begin{aligned} \log p_\theta(x, y) &= \log p_\theta(y | x) + \log p_\theta(z) + \log \frac{p_\theta(x | z)}{p_\theta(z | x)} \\ &= \int q_\phi(z | x, y) \left\{ \log p_\theta(y | x) - \log \frac{q_\phi(z | x, y)}{p_\theta(z)} + \right. \\ &\quad \left. \log \frac{q_\phi(z | x, y) p_\theta(x | z)}{p_\theta(z | x)} \right\} dz \\ &\geq \mathbb{E}_{q_\phi(z|x, y)} [\log p_\theta(y | x)] - D_{KL} [q_\phi(z | x, y) \| p_\theta(z)], \end{aligned} \quad (4)$$

where ϕ and θ are modeled by neural networks. The first term in Eq. (4) aims to maximize the likelihood by improving the confidence of the prediction. When optimizing the first term, the optimizing process can be written as

$$\begin{aligned} &\arg \max_{q_\phi(z|x, y)} \int_x \sum_y p(x, y) \int_z q_\phi(z | x, y) \log p_\theta(y | x) \\ &= \arg \max_{q_\phi(z|x, y)} \int_x p(x) \int_z q_\phi(z | x, y) \sum_y p(y | x) \log p_\theta(y | x). \end{aligned} \quad (5)$$

As $q_\phi(z | x, y)$ is a probability distribution, the integral is upper bounded by $\max_y \sum_y p(y | x) \log p_\theta(y | x)$, which is equivalent to minimizing the cross entropy loss in Eq. (1). The second KL divergence minimizes the divergence between the variational posterior distribution $q_\phi(z|x)$ and the

prior distribution $p_\theta(z)$. The third term actually can be calculated by $D_{KL} [q_\phi(z | x, y) \| p_\theta(z | x) / p_\theta(x | z)]$, which is a non-negative value and we eliminate this item. In order to explicitly contain the covariance of all distributions, the remaining problem lies in how to formulate the posterior distribution $q_\phi(z|x)$ and the prior distribution $p_\theta(z)$.

We first assume that z satisfies multivariate Gaussian with a mean and diagonal covariance. The $p_\theta(z)$ can be formulated as $\mathcal{N}(\mu_c, \mathbf{I})$ with mean μ_c as the class center and fixed eye matrix \mathbf{I} , respectively. As we aim to constrain the covariance, we adopt the fixed covariance as the prior. In order to learn the latent distribution quickly from a single instance during training, we adopt linear layers $P_\mu(\cdot)$ and $P_\sigma(\cdot)$ after the feature extraction network to form a statics prediction pipeline to estimate $q_\phi(z|x, y)$ following previous works [16]

$$\begin{aligned} q_\phi(z|x, y) &= \mathcal{N}(z | \hat{\mu}_c, \hat{\sigma}_c^2) \\ &= [P_\mu(x), P_\sigma(x)]_{y=c}. \end{aligned} \quad (6)$$

After obtaining the estimated statistical information, we can then formulate the second KL divergence term into the covariance constraint loss by

$$\begin{aligned} &-D_{KL} [q_\phi(z | x, y) \| p_\theta(z)] \\ &= \int q_\phi(z|x, y) \log \frac{p_\theta(z)}{q_\phi(z|x, y)} dz \\ &= \int \mathcal{N}(z; \hat{\mu}, \hat{\sigma}^2) \log \frac{\mathcal{N}(z; \mu_c, \mathbf{I})}{\mathcal{N}(z; \hat{\mu}_c, \hat{\sigma}_c^2)} dz \\ &= \frac{1}{2} \sum_{c=1}^{c^0} \sum_{i=1}^d (1 + \log \hat{\sigma}_{c,i}^2 - \hat{\sigma}_{c,i}^2 - (\hat{\mu}_c - \mu_c)^2) \\ &\leq \frac{1}{2} \sum_{c=1}^{c^0} \sum_{i=1}^d (1 + \log \hat{\sigma}_{c,i}^2 - \hat{\sigma}_{c,i}^2), \end{aligned} \quad (7)$$

where d represents the feature dimension. We drop the final term to formulate the covariance constraint loss as $\mathcal{L}_{ccl} = \frac{1}{2} \sum_{c=1}^{c^0} \sum_{i=1}^d (1 + \log \hat{\sigma}_{c,i}^2 - \hat{\sigma}_{c,i}^2)$, which constrains the covariance of each feature distribution. Eq. (7) shows that the covariance constraint loss can be viewed as the upper bound of $-D_{KL}$. Thus, by deploying covariance constraint loss, we are able to obtain a tighter upper bound than Eq. (4), which can be easier to optimize. The overall learning objective for the base session training is

$$\mathcal{L}_{Base} = \mathcal{L}_{ce} + \gamma \mathcal{L}_{ccl}, \quad (8)$$

where γ is a positive hyperparameter that controls the degree of constraint.

3.4. Semantic Perturbation Learning

In order to prevent the model from overfitting the few-shot new classes during incremental sessions, recent approaches [10, 18, 29] attempt to freeze the feature extractor

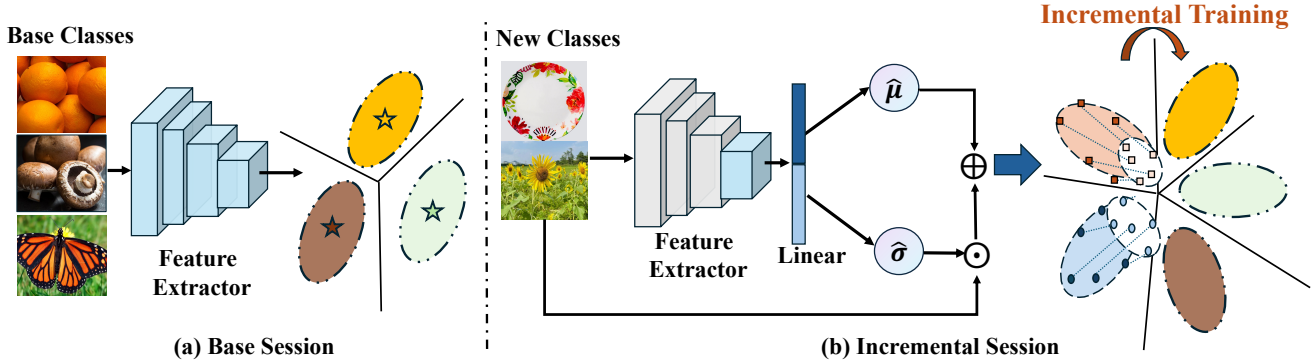


Figure 2. (a) Base session training. We deploy the covariance constraint loss to learn class distributions with a fixed span. (b) Semantic perturbation learning for the incremental stages. New data samples are perturbed by multiplying with the predicted distribution. The perturbed samples are trained along with the original data to establish the separation between classes.

and learn the new class prototypes by averaging the few-shot samples. Though the overfitting problem is mitigated by fixing the previously learned feature space, the new class distributions can only account for small parts of the feature space compared to base classes due to the scarcity of data. Under this ill-divided space dilemma, as shown in Fig. 1(b), new class prototypes may lie very close to base class distributions or new class prototypes, which may further lead the model to misclassify new classes.

In this paper, we propose to address this issue from two perspectives. Firstly, we intend to expand the new class feature distributions by assigning the fixed covariance, which is the same as the base distributions and allows the new class distributions to take up more feature space. Secondly, to facilitate better separation among classes, we push the few-shot samples to be away from semantically similar distributions and then retrain the classifier to distinguish the perturbed new class samples from other classes. We achieve these two goals by forming the semantic perturbation learning framework, where we aim to learn a perturbation distribution from which any perturbations can perturb new data to be away from those close distributions but within a fixed range. The perturbation distribution can also be learned by formulating a similar variational inference in Eq. (4) but with a different prior distribution. Specifically, we form the prior distribution as the Gaussian distribution $p_\theta(\tilde{z})$ as $\mathcal{N}(\tilde{\mu}, \mathbf{I})$, where $\tilde{\mu}$ is the linear combination of other class prototypes. The weight is calculated by the similarity score over classes. For each incremental session, we first initialize the classifier by computing the new class prototypes by averaging the few-shot samples. Given new data sample x_i^t with label y_i^t , we compute the similarity score over other classes:

$$S_{i,c} = \frac{\mathbb{I}_{c \neq y_i^t} \exp(\cos(\mathbf{w}_c, x_i^t))}{\sum_{c=1}^{|\mathcal{C}^t|-1} \exp(\cos(\mathbf{w}_c, x_i^t))}, \quad (9)$$

where $S \in \mathbb{R}^{K \times (|\mathcal{C}^t|-1)}$ and $\mathbb{I}(\cdot)$ is an indicator function. For each sample, we can get the prior mean by multiplying the confidence score with other class prototypes

$$\tilde{\mu} = \sum_{c \neq y_i^t}^{|\mathcal{C}^t|-1} S_{i,c} \mathbf{w}_c. \quad (10)$$

We also adopt a linear layer to predict the mean and covariance of each sample, the same as Eq. (6). To be noted, we do not reuse the linear layer in the base session so that this method can be directly applied to other methods in a plug-and-play manner. After predicting the statistics, the overall objective during the incremental training is:

$$\mathcal{L}_{inere} = \mathcal{L}_{ce}(x_i^t, y_i^t) + \mathcal{L}_{ce}(\hat{\mu} + \hat{\sigma} \odot x_i^t, y_i^t) - \alpha D_{KL}[q_\phi(\tilde{z} | x_i^t) \| p_\theta(\tilde{z} | x_i^t)], \quad (11)$$

where \odot represents the element-wise multiplication. The first and second term is the classification loss for the few-shot data and the perturbed few-shot data. The third term aims to learn the perturbation distribution by minimizing the KL divergence between the predicted feature distribution of each sample and the prior distribution. The positive hyperparameter α controls the strength of the perturbation. The KL term is formulated as:

$$\begin{aligned} & -D_{KL}[q_\phi(\tilde{z} | x_i^t) \| p_\theta(\tilde{z} | x_i^t)] \\ & = \frac{1}{2} \sum_{i=1}^d (1 + \log \hat{\sigma}_i^2 - \hat{\sigma}_i^2 - (\hat{\mu}_i - \tilde{\mu}_i)^2). \end{aligned} \quad (12)$$

4. Experiments

4.1. Benchmark Performance

We conduct our experiments on three FSCIL benchmarks, i.e., CIFAR100, MiniImageNet, and CUB200, and compare our approach with the baseline models and other

Table 1. Incremental learning performance on MiniImageNet under 5-way 5-shot setup. “Avg Acc.” represents the average accuracy of all sessions. “Final Improv.” calculates the improvement of our method after learning in the final session. **Bold** represents best performance. * indicates that we reproduce the results using public open-source code.

Methods	Accuracy in each session (%) \uparrow										Avg Acc.	Final Improv.
	0	1	2	3	4	5	6	7	8			
iCaRL [26]	61.31	46.32	42.94	37.63	30.49	24.00	20.89	18.80	17.21	33.29	+39.20	
NCM [11]	61.31	47.80	39.30	31.90	25.70	21.40	18.70	17.20	14.17	30.83	+42.24	
Data-free Replay [22]	71.84	67.12	63.21	59.77	57.01	53.95	51.55	49.52	48.21	58.02	+8.20	
CEC [18]	72.00	66.83	62.97	59.43	56.70	53.73	51.19	49.24	47.63	57.75	+8.78	
MetaFSCIL [5]	72.04	67.94	63.77	60.29	57.58	55.16	52.90	50.79	49.19	58.85	+7.22	
C-FSCIL [10]	76.40	71.14	66.46	63.29	60.42	57.46	54.78	53.11	51.41	61.61	+5.00	
LIMIT [51]	73.81	72.09	67.87	63.89	60.70	57.77	55.67	53.52	51.23	61.84	+5.18	
CE	75.65	70.45	66.09	62.16	58.96	55.92	53.08	51.05	49.39	60.31	+7.02	
CE-Ours	75.77	70.71	66.53	63.33	60.42	57.46	54.61	52.48	50.65	61.33	+5.76	
FACT* [49]	76.25	70.91	66.41	62.79	59.45	56.22	53.37	51.21	49.48	60.68	+6.93	
FACT*-Ours	75.65	70.63	66.79	63.20	59.72	56.68	53.78	51.53	50.01	60.89	+6.40	
SAVC* [29]	80.68	75.89	71.54	67.80	64.85	61.42	58.38	56.43	54.91	65.77	+1.50	
SAVC*-Ours	80.90	75.89	71.80	68.59	65.86	62.41	59.33	57.71	56.41	66.54		




Figure 3. (a) Comparison of harmonic performance after incremental sessions with baseline models on MiniImageNet. (b) Comparison of performance of new classes with baseline model in each incremental session on MiniImageNet

recent FSCIL methods. We also compare our method with classical CIL methods such as ICaRL [26], NCM [11], and FSCIL methods CEC [18], C-FSCIL [10] and meta-learning based approaches [4, 51]. We report the numerical results of MiniImageNet in Tab. 1, and more results on CIFAR100 and CUB200 can be found in Appendix C.

For baseline models, adding our proposed loss and semantic perturbation learning boosts the model performance for incremental sessions. After 8-session incremental learning, our method can improve the final model performance by at least 0.53% on MiniImageNet and 0.56% on CIFAR100, respectively. It is worth noticed that our method has a more profound effect on SAVC, which learns representations by pushing the data to the class center using contrastive learning. As SAVC does not consider the span of each class distribution explicitly, constraining the class dis-

tribution using our proposed methods boosts the base class performance and obtains better performance when learning new classes.

We further demonstrate the comparisons of harmonic performances of all classes after incremental sessions and comparisons of new class performance in each incremental session in Fig. 3. As shown in Fig. 3(a), compared with baseline models, our approach boosts the harmonic performance by at least 9.80%. The CE model enjoys an improvement of 11.45% from our approach. The significant boost in harmonic performance results from the improvement of new classes. As illustrated in Fig. 3(b), the performance of new classes demonstrates a substantial improvement at each incremental session. Overall, our approach boosts the new class performance by at least 5.65%. Benefiting from the covariance constraint and distribution expansion strat-

egy, we are able to learn more separable representations of new classes, boosting the performance of new classes.

4.2. Implementation Details

Datasets and Experimental Settings. We conduct extensive experiments on three benchmark datasets, including MiniImageNet, CIFAR100 and CUB200. MiniImageNet and CIFAR100 contain 100 classes in total, we set the number of base classes as 60. We set 8 incremental sessions, where each session formulates a 5-way 5-shot problem. For fine-grained data CUB200 that contains 200 classes, we set the number of base classes as 100, followed by 10 incremental sessions, and each session formulates a 10-way 5-shot problem. To make the comparison fair, we use the same base and incremental class data in each dataset when conducting the experiments following previous works [1, 29, 49, 51]. All experiments are conducted using one RTX3090 card¹.

Baseline Models. We adopt three baseline models, i.e., the naive CE-trained model, fantasy-based model FACT [49], and state-of-the-art contrastive learning model SAVC [29]. Details of the baseline models can be found in Appendix A.

Model Architectures. Following [18, 29, 38, 51], we use ResNet-18 in experiments on CIFAR100 and MiniImageNet. We use ImageNet pre-trained ResNet-18 for the CUB200 dataset following [29, 38, 51]. We follow the same experimental setting as baseline methods for fair comparisons. The dimension d is set as 64, 512, 512 on CIFAR100, MiniImageNet and CUB200, respectively, which is the same as baseline models. The hyperparameter analysis is shown in Sec. 4.5.2.

4.3. Ablation Studies of Each Component

We conduct the ablation experiments of the proposed CCL and SPL on the CIFAR100 dataset, as shown in Tab. 2. Following previous works [29, 38, 49], we verify the effectiveness of our approach by six metrics, i.e., the performance of the base model, the performance on the base classes after incremental sessions, the performance of all new classes, the average performance of all incremental sessions, the performance drop of the base classes ($PD = Accuracy_{base} - Accuracy_{old}$), and the harmonic accuracy of base and new classes after the incremental learning. When we do not deploy the SPL for incremental sessions, the model adopts the prototype-based incremental learning approach mentioned in Sec. 3.1.2.

As shown in Tab. 2, by deploying our proposed covariance constraint loss L_{ccl} , our model is able to achieve higher base class performance compared to the baseline model. By constraining the base class distributions, the drop in PD rate and the gain of new class performance on three baseline

Table 2. Ablation results of proposed components. “Base” represents the base model performance. “Old” is the performance of base classes after incremental sessions. “New” is the performance of all new classes. “Avg” is the average performance of all incremental sessions. “PD” is the drop rate of the base class performance after incremental sessions. “H.” is the harmonic accuracy of base and new classes after incremental learning.

Method	Base Session	Incremental Session	CIFAR100					
	L_{ccl}	SPL	Base \uparrow	Old \uparrow	New \uparrow	Avg \uparrow	PD \downarrow	H. \uparrow
CE			76.87	71.15	20.65	62.12	5.72	32.01
CE+SPL		✓	76.87	69.32	23.38	62.20	7.55	34.96
CE+ L_{ccl}	✓		78.26	72.67	21.90	63.58	5.59	33.66
CE+ L_{ccl} +SPL	✓	✓	78.26	71.42	23.62	63.66	6.84	35.51
FACT			78.38	71.08	21.73	62.17	5.40	32.54
FACT+SPL		✓	78.38	71.07	21.82	62.20	7.31	33.40
FACT+ L_{ccl}	✓		79.12	72.17	22.50	62.80	4.73	34.31
FACT+ L_{ccl} +SPL	✓	✓	79.12	71.20	24.05	62.82	5.70	35.96
SAVC			78.98	72.25	21.00	62.81	6.73	32.56
SAVC+SPL		✓	78.98	72.20	21.25	62.88	6.78	32.83
SAVC+ L_{ccl}	✓		79.00	72.88	22.38	63.35	6.12	34.24
SAVC+ L_{ccl} +SPL	✓	✓	79.00	72.90	23.50	63.48	6.10	35.51

models validate that the covariance constraint loss allows the model to learn the new classes better while obtaining less forgetting of old classes.

We also conduct experiments of SPL on both baseline models and L_{ccl} trained models. The results demonstrate that directly deploying SPL can boost the performance of new classes and harmonic accuracy by at least 0.25% and 0.27%, respectively. It is worth noticing that simply using SPL on the baseline models leads to a higher PD. As we expand the distribution of new classes during incremental sessions without constraining the base class distributions, it may lead the model to confuse the old classes with new classes. If SPL is deployed on the L_{ccl} trained model, the model is able to achieve the highest performance on new classes and the highest harmonic accuracy, validating the effectiveness of the combination of proposed components. We also conduct experiments on comparing SPL with other incremental update methods in Appendix B.

4.4. T-SNE Visualizations of Proposed Approaches

We compare the different approaches by visualizing the learned feature embedding on the CIFAR100 dataset. In Fig. 4(a), we compare the base model trained with and without covariance constraint loss. We visualize the feature embedding of 8 different base classes, demonstrating that by adding the constraint on the covariance of the feature distribution, the learned feature distribution becomes more separable and compact, leaving more feature space for the incoming new classes. In Fig. 4(b), we demonstrate the effectiveness of SPL by showing the learned new class features before SPL, and new class features along with all perturbed data after SPL during incremental sessions ($t > 0$). As shown in the left side of Fig. 4(b), the new class features (colored triangles) tend to lie close with each other

¹<https://github.com/tambourine666/Covariance-Space-Regularize>

and mix with base classes due to the scarcity of the training data. By deploying SPL during the incremental sessions, our approach is able to push the few-shot training samples away from base classes and facilitate better separation among classes. The generated perturbed data in each epoch (denoted by “×” in the figure) expands the original small new class feature distribution, which allows the new class distributions to take more feature space, benefiting the new class prediction.

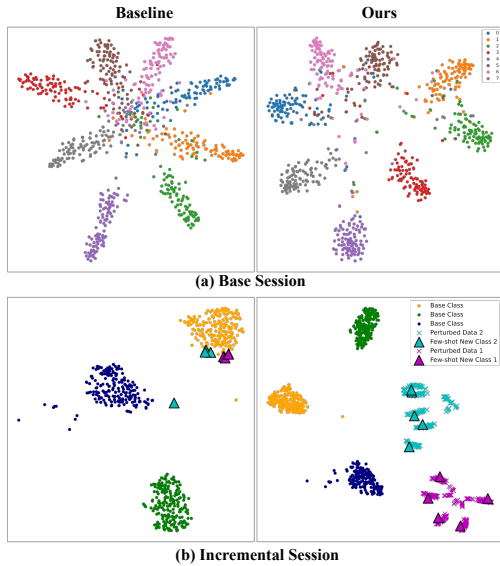


Figure 4. Comparisons of T-SNE visualization of the learned feature embedding on CIFAR100. (a) Visualization of eight base classes’ feature embedding in the base session. We compare the base model with and without covariance constraint loss. (b) Visualizations of the feature embedding of two few-shot new classes together with their perturbed samples, and three base classes during incremental learning. We compare the baseline model and the model using the proposed SPL.

4.5. Further Analysis

4.5.1 Effects of Incremental Shots

We further conduct experiments on CUB200 to investigate the impact of varying the number of shots during incremental sessions, as shown in Fig. 5(a). Since our method relies on the data of each class to expand the new class distributions, we vary the number of shots per class to observe its effect in incremental sessions. Keeping the number of classes consistent with the current experimental setting, we vary the shot number K from $\{1, 3, 5, 10\}$ on CUB200. As depicted in the figure, increasing the number of available instances in each class leads to more accurate distributions and a corresponding improvement on performance. Even with a reduced shot number of 1, the model maintains stable performance during incremental sessions.

4.5.2 Hyperparameter Analysis

There are two hyperparameters in our approach, i.e., the γ controls the impact of covariance constraint loss in Eq. (8) and the α controls the perturbation in Eq. (11). We conduct experiments by changing γ and α from $\{0, 0.0001, 0.001, 0.01, 0.1\}$, and show the accuracy of the last session on CUB200 dataset in Fig. 5 (b). We can see that relatively small γ and α yield better performance, and the best performance is yield with γ and α set as 0.01.

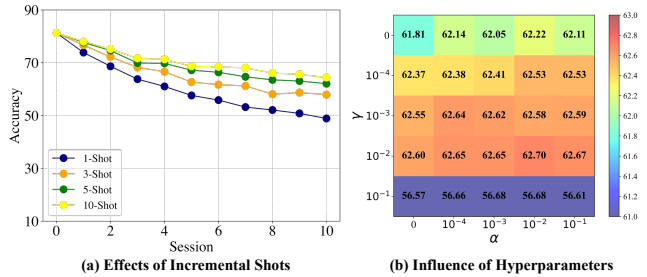


Figure 5. (a) Results of different incremental shots on CUB200. (b) Results of different hyperparameters on CUB200.

5. Conclusion

In this paper, we propose a covariance constraint loss and semantic perturbation learning to address the ill-divided feature space problem for few-shot class incremental learning (FSCIL). Our main motivation is to constrain the span of each distribution then reallocate the ill-divided feature space by reestablishing the decision boundaries between classes. Based on this, we attempt to learn the model through two steps: during the base session, we propose a covariant constraint loss (CCL) to explicitly constrain the feature distribution and facilitate better class separation. We derive the CCL from a variational inference framework, which estimates and constrains the feature distribution efficiently. For incremental sessions, we propose to generate semantic-guided perturbed data to aid the learning of few-shot new classes. The generated data expands the few-shot distributions and pushes the few-shot samples away from easily confusing classes. The proposed approach can be integrated into current FSCIL methods in a plug-and-play manner, which is easy to implement. We conduct comprehensive experiments on three benchmark datasets and apply our approach to three baseline models. Experimental results demonstrate the effectiveness of our methods and obtain a new state-of-the-art performance.

6. Acknowledgements

The work was partially supported by the following: National Natural Science Foundation of China under no. 92370119 and No. 62276258, and No.62376113.

References

- [1] Aishwarya Agarwal, Biplab Banerjee, Fabio Cuzzolin, and Subhasis Chaudhuri. Semantics-driven generative replay for few-shot class incremental learning. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5246–5254, 2022. [7](#)
- [2] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 233–248, 2018. [3](#)
- [3] Xiaohua Chen, Yucan Zhou, Dayan Wu, Wanqian Zhang, Yu Zhou, Bo Li, and Weiping Wang. Imagine by reasoning: A reasoning-based implicit semantic data augmentation for long-tailed classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 356–364, 2022. [3](#)
- [4] Meng Cheng, Hanli Wang, and Yu Long. Meta-learning-based incremental few-shot object detection. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(4):2158–2169, 2021. [6](#)
- [5] Zhixiang Chi, Li Gu, Huan Liu, Yang Wang, Yuanhao Yu, and Jin Tang. Metafscl: A meta-learning approach for few-shot class incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14166–14175, 2022. [4](#), [6](#), [13](#)
- [6] Shehzaad Dhuliawala, Mrinmaya Sachan, and Carl Allen. Variational classification. *TMLR*, 2023. [4](#)
- [7] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013. [1](#)
- [8] Dandan Guo, Long Tian, He Zhao, Mingyuan Zhou, and Hongyuan Zha. Adaptive distribution calibration for few-shot learning with hierarchical optimal transport. *Advances in Neural Information Processing Systems*, 35:6996–7010, 2022. [3](#)
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [1](#)
- [10] Michael Hersche, Geethan Karunaratne, Giovanni Cherubini, Luca Benini, Abu Sebastian, and Abbas Rahimi. Constrained few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9057–9067, 2022. [2](#), [3](#), [4](#), [6](#), [13](#)
- [11] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 831–839, 2019. [3](#), [6](#), [13](#)
- [12] Minsoo Kang, Jaeyoo Park, and Bohyung Han. Class-incremental learning by knowledge distillation with adaptive feature consolidation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16071–16080, 2022. [2](#)
- [13] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. [12](#)
- [14] Do-Yeon Kim, Dong-Jun Han, Jun Seo, and Jaekyun Moon. Warping the space: Weight space rotation for class-incremental few-shot learning. In *The Eleventh International Conference on Learning Representations*, 2022. [3](#)
- [15] Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27, 2014. [2](#), [4](#)
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#), [4](#)
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. [1](#)
- [18] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion*, 58:52–68, 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [13](#)
- [19] Yushu Li, Xun Xu, Yongyi Su, and Kui Jia. On the robustness of open-world test-time training: Self-training with dynamic prototype expansion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11836–11846, 2023. [3](#)
- [20] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. [2](#), [3](#)
- [21] Binghao Liu, Boyu Yang, Lingxi Xie, Ren Wang, Qi Tian, and Qixiang Ye. Learnable distribution calibration for few-shot class-incremental learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [3](#), [13](#)
- [22] Huan Liu, Li Gu, Zhixiang Chi, Yang Wang, Yuanhao Yu, Jun Chen, and Jin Tang. Few-shot class-incremental learning via entropy-regularized data-free replay. In *European Conference on Computer Vision*, pages 146–162. Springer, 2022. [6](#), [13](#)
- [23] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost Van De Weijer. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5513–5533, 2022. [2](#)
- [24] 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](#)
- [25] Can Peng, Kun Zhao, Tianren Wang, Meng Li, and Brian C Lovell. Few-shot class-incremental learning from an open-set perspective. In *European Conference on Computer Vision*, pages 382–397. Springer, 2022. [2](#)
- [26] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. [2](#), [6](#), [13](#)

- [27] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [28] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017. [2, 3](#)
- [29] Zeyin Song, Yifan Zhao, Yujun Shi, Peixi Peng, Li Yuan, and Yonghong Tian. Learning with fantasy: Semantic-aware virtual contrastive constraint for few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24183–24192, 2023. [2, 3, 4, 6, 7, 12, 13](#)
- [30] Zhaorui Tan, Xi Yang, and Kaizhu Huang. Rethinking multi-domain generalization with a general learning objective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23512–23522, 2024. [1](#)
- [31] Zhaorui Tan, Xi Yang, Qifeng Wang, Anh Nguyen, and Kaizhu Huang. Interpret your decision: Logical reasoning regularization for generalization in visual classification. *arXiv preprint arXiv:2410.04492*, 2024. [1](#)
- [32] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12183–12192, 2020. [1, 3](#)
- [33] Songsong Tian, Lusi Li, Weijun Li, Hang Ran, Xin Ning, and Prayag Tiwari. A survey on few-shot class-incremental learning. *arXiv preprint arXiv:2304.08130*, 2023. [1](#)
- [34] Yonglong Tian, Yue Wang, Dilip Krishnan, Joshua B Tenenbaum, and Phillip Isola. Rethinking few-shot image classification: a good embedding is all you need? In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 266–282. Springer, 2020. [3](#)
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#)
- [36] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *International conference on machine learning*, pages 6438–6447. PMLR, 2019. [12](#)
- [37] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016. [3](#)
- [38] Qi-Wei Wang, Da-Wei Zhou, Yi-Kai Zhang, De-Chuan Zhan, and Han-Jia Ye. Few-shot class-incremental learning via training-free prototype calibration. *Advances in Neural Information Processing Systems*, 36, 2024. [2, 7](#)
- [39] Yulin Wang, Xuran Pan, Shiji Song, Hong Zhang, Gao Huang, and Cheng Wu. Implicit semantic data augmentation for deep networks. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [40] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 374–382, 2019. [3](#)
- [41] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3014–3023, 2021. [3](#)
- [42] Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, Qing Yang, and Cheng-Lin Liu. Convolutional prototype network for open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(5):2358–2370, 2020. [2, 3](#)
- [43] Shuo Yang, Songhua Wu, Tongliang Liu, and Min Xu. Bridging the gap between few-shot and many-shot learning via distribution calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9830–9843, 2021. [3](#)
- [44] Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning. *arXiv preprint arXiv:2302.03004*, 2023. [2, 3, 12](#)
- [45] Jaehong Yoon, Sultan Madjid, Sung Ju Hwang, Chang-Dong Yoo, et al. On the soft-subnetwork for few-shot class incremental learning. In *International Conference on Learning Representations (ICLR) 2023*. International Conference on Learning Representations, 2023. [3](#)
- [46] Tianyuan Yu, Sen He, Yi-Zhe Song, and Tao Xiang. Hybrid graph neural networks for few-shot learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 3179–3187, 2022. [3](#)
- [47] Chi Zhang, Nan Song, Guosheng Lin, Yun Zheng, Pan Pan, and Yinghui Xu. Few-shot incremental learning with continually evolved classifiers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12455–12464, 2021. [1, 3](#)
- [48] Jian Zhang, Chenglong Zhao, Bingbing Ni, Minghao Xu, and Xiaokang Yang. Variational few-shot learning. In *Proceedings of the IEEE/CVF International conference on computer vision*, pages 1685–1694, 2019. [2](#)
- [49] Da-Wei Zhou, Fu-Yun Wang, Han-Jia Ye, Liang Ma, Shiliang Pu, and De-Chuan Zhan. Forward compatible few-shot class-incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9046–9056, 2022. [2, 3, 4, 6, 7, 12, 13](#)
- [50] Da-Wei Zhou, Qi-Wei Wang, Han-Jia Ye, and De-Chuan Zhan. A model or 603 exemplars: Towards memory-efficient class-incremental learning. In *The Eleventh International Conference on Learning Representations*, 2022. [2](#)
- [51] Da-Wei Zhou, Han-Jia Ye, Liang Ma, Di Xie, Shiliang Pu, and De-Chuan Zhan. Few-shot class-incremental learning by sampling multi-phase tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [3, 6, 7, 13](#)
- [52] Fei Zhu, Zhen Cheng, Xu-yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. *Advances in Neural Information Processing Systems*, 34:14306–14318, 2021. [1, 2](#)

- [53] Fei Zhu, Xu-Yao Zhang, Chuang Wang, Fei Yin, and Cheng-Lin Liu. Prototype augmentation and self-supervision for incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5871–5880, 2021. [3](#)
- [54] Kai Zhu, Yang Cao, Wei Zhai, Jie Cheng, and Zheng-Jun Zha. Self-promoted prototype refinement for few-shot class-incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6801–6810, 2021. [2](#), [3](#), [13](#)
- [55] Yixiong Zou, Shanghang Zhang, Yuhua Li, and Ruixuan Li. Margin-based few-shot class-incremental learning with class-level overfitting mitigation. *Advances in neural information processing systems*, 35:27267–27279, 2022. [4](#)