

Learning with Fantasy: Semantic-Aware Virtual Contrastive Constraint for Few-Shot Class-Incremental Learning

Zeyin Song^{1†} Yifan Zhao^{2†} Yujun Shi³ Peixi Peng^{2,4*} Li Yuan^{1,4} Yonghong Tian^{1,2,4*}

¹School of Electronic and Computer Engineering, Peking University

²School of Computer Science, Peking University

³National University of Singapore

⁴Peng Cheng Laboratory

zysong@stu.pku.edu.cn, {zhaoyf, pxpeng, yuanli-ece, yhtian}@pku.edu.cn, shi.yujun@u.nus.edu


Abstract

Few-shot class-incremental learning (FSCIL) aims at learning to classify new classes continually from limited samples without forgetting the old classes. The mainstream framework tackling FSCIL is first to adopt the cross-entropy (CE) loss for training at the base session, then freeze the feature extractor to adapt to new classes. However, in this work, we find that the CE loss is not ideal for the base session training as it suffers poor class separation in terms of representations, which further degrades generalization to novel classes. One tempting method to mitigate this problem is to apply an additional naïve supervised contrastive learning (SCL) in the base session. Unfortunately, we find that although SCL can create a slightly better representation separation among different base classes, it still struggles to separate base classes and new classes. Inspired by the observations made, we propose Semantic-Aware Virtual Contrastive model (SAVC), a novel method that facilitates separation between new classes and base classes by introducing virtual classes to SCL. These virtual classes, which are generated via pre-defined transformations, not only act as placeholders for unseen classes in the representation space, but also provide diverse semantic information. By learning to recognize and contrast in the fantasy space fostered by virtual classes, our SAVC significantly boosts base class separation and novel class generalization, achieving new state-of-the-art performance on the three widely-used FSCIL benchmark datasets. Code is available at: <https://github.com/zysong0113/SAVC>.

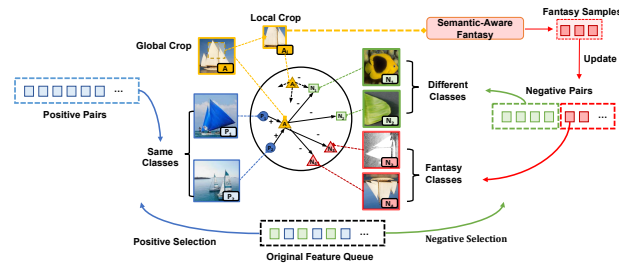
1. Introduction

Like humans, the modern artificial intelligence models are expected to be equipped with the ability of learning continually. For instance, a face recognition system needs

[†]: authors contributed equally. *: corresponding author.



(a) Comparison of the embedding spaces trained with different methods.



(b) Illustration of our Semantic-Aware Virtual Contrastive model (SAVC).

Figure 1. The motivation of the proposed approach. Under the incremental-frozen framework in FSCIL, our SAVC learns to recognize and contrast in the fantasy space, which leads to better base class separation and novel class generalization than CE and SCL.

to continually learn to recognize new faces without forgetting those has already kept in mind. In this case, Class-Incremental Learning (CIL) [19, 25, 30, 37, 50] draws attention extensively, that sufficient and disjoint data are coming session-by-session. However, it is unrealistic to always rely on accessing of multitudinous data. Taking the case of face recognition system again, whether it can quickly learn a new concept with just a few images, is an important criterion for evaluating its performance. As a result, designing effective and efficient algorithms to resolve Few-Shot Class-Incremental Learning (FSCIL) problem has drawn increasing attention [6, 8, 12, 40, 44, 51, 55]. Unlike CIL, only a few data are provided for FSCIL at incremental session

while sufficient samples are only visible at the base session.

The core challenge of CIL is to strike a balance of stability and plasticity, *i.e.*, suppressing forgetting on old knowledge while adapting smoothly to new knowledge. Previous works on FSCIL have demonstrated that an incrementally trained model is prone to overfit on limited new data and suffer catastrophic forgetting [8, 12, 44, 54]. Thus the mainstream FSCIL framework [40, 51, 58] is to first use the cross-entropy (CE) loss for training in the base session, then freeze the backbone to adapt to new classes. Given such baseline framework, the main problem that needs to be taken into consideration is: *What makes a good base-session-model that can generalize well to new classes with limited data?*

Intuitively, if all representations of base classes concentrate around their own clustering centers (*i.e.*, prototypes), and all prototypes are far away from each other, then novel classes should be easily incorporated into the current space without overlapping. Therefore, we conjecture that base class separation facilitates novel class generalization. However, the current widely adopted CE loss in FSCIL cannot well separate the class margins, which causes poor generalization to new classes. This motivates our attempt at methods with better clustering effects, such as supervised contrastive learning (SCL) [21]. We conduct experiments to validate our conjecture by training models optimized by only CE loss, and both CE and SCL. We empirically find that CE leads to lower base class separation and classification accuracy than SCL. Although SCL indeed slightly improves the base class separation with higher accuracy, we find it implicitly decreases the inter-class distance. Therefore, the two methods both leave inadequate room for future updates, which leads to the inevitable overlapping between novel classes and old classes in the representation space (as shown in the left and middle columns of Fig. 1a).

The unsatisfactory performance of SCL implies that the limited semantic information base session provided may need our fantasy remedy. Although we have no access to future classes in the base session, we can create many virtual classes with various semantics filling the unallocated representation space, which not only holds enough space for future classes, but also enables the learning of diverse semantics for better inference and generalization.

To this end, we propose a Semantic-Aware Virtual Contrastive (SAVC) model to improve the FSCIL by boosting base class separation (see Fig. 1b). Our fantasy space is constructed with pre-defined transformations [11, 15, 35, 52], which can be regarded as the semantic extension of the original space in finer grains. The type of transformations represents the richness of our imagination, *i.e.*, how compact we want the embedding space to be and how fine the semantic grains we want to acquire. Learning to contrast and recognize in the fantasy space not only brings a better class

separation effect, but also provides richer semantic information to improve the accuracy. In contrast to the CE baseline and SCL, our SAVC can simultaneously decrease intra-class distance and increase inter-class distance, which leads to better base class separation and novel class generalization (as given in the right of Fig. 1a). Extensive experiments on various benchmark datasets further verify the superiority of our method.

Our main contributions are summarized in three folds:

- 1) We empirically discover that it is crucial to boost class separation degree in base session for the incremental-frozen framework in FSCIL, which helps fast generalization for novel classes with only a few samples.

- 2) We propose a novel Semantic-Aware Virtual Contrastive (SAVC) framework, enhancing the FSCIL learner by learning to recognize and contrast in the fantasy space, which preserves enough place for novel generalization and enables a multi-semantic aggregated inference effect.

- 3) Extensive experiments on three FSCIL benchmarks, *i.e.*, CIFAR100, *mini*ImageNet and CUB200, demonstrate that our SAVC outperforms all approaches by a large margin and achieves the state-of-the-art performance.

2. Related Work

Few-Shot Learning. FSL focuses on fast adaptation to unseen classes through few learnable samples. Current FSL methods are mainly divided into two lines: optimization-based methods [14, 26, 33, 36] and metric-based methods [34, 42, 43, 45]. Optimization-based methods aim to encourage the model to obtain generalizable representations by using meta-learning frameworks. Our method has more connection with metric-based methods, which use a pre-trained backbone to extract feature, and compute score by measuring the distance between query and support. Nevertheless, FSL aims to adapt to novel classes via few samples and cannot be directly applied to FSCIL, without maintaining classification performance over base classes.

Class Incremental Learning. CIL is the process of learning new knowledge continually without forgetting. Classical works fall into three streams. The first stream rehearses former samples selectively to maintain discrimination over old classes [5, 37, 38]. The second stream regularizes the important parameters from being dramatically changed [25, 29]. The third stream dynamically expands the network architecture to adjust the model’s representation ability on new tasks [49, 50]. Recent methods boost performance by encouraging the model to learn to recognize more classes [41, 57], which are relative to our method. CIL methods build on the basis of the large number of training examples, which is unsuitable for FSCIL.

Few-Shot Class-Incremental Learning. FSCIL [1, 2, 6, 9, 10, 18, 27, 31] aims at undertaking CIL task in insufficient data settings. TOPIC [44] first proposes this new task

and solves it by neural gas structure. Current FSCIL methods can be roughly divided into two groups. The first group updates the backbone in incremental session [8, 12, 44, 54], while the second group fixes the backbone in incremental sessions [40, 51, 58]. Built on the latter incremental-frozen framework, other works balance the stability and plasticity trade-off from different perspectives, such as finding flatten minima [40], episodic training [51, 56, 58], ensemble learning [20], forward compatibility [55], and so on.

Self-supervised learning. SSL aims at learning representations without human annotations. A common way of SSL is to design pretext tasks and teach the model how to solve these prediction tasks [11, 15, 35, 52]. [24] proposes to utilize these transformations under full-supervised datasets and trains a unified classifier, which motivates the construction of our fantasy space. Another commonly used pretext task is instance discrimination represented by contrastive learning [3, 7, 16]. The methodology is to teach the model to aggregate positives and exclude negatives. SupCon [21] converted to “class discrimination” task, leveraging the label information to decide positive pairs.

3. Methodology

3.1. Preliminary


In FSCIL, we learn a sequence of tasks continually. There is a constant stream of training data denoting $\mathcal{D}_{train} = \{\mathcal{D}_{train}^t\}_{t=0}^T$, where $\mathcal{D}_{train}^t = \{(\mathbf{x}_i, y_i)\}_{i=0}^{N_t}$ represents the training samples from session t , \mathbf{x}_i and y_i are the i -th image and corresponding label respectively. Assuming that the label space of the t -th task is \mathcal{C}^t , the label spaces from different tasks have no overlap, *i.e.*, the model trained on the \mathcal{D}_{train}^t should be evaluated on \mathcal{D}_{test}^t which contains all encountered classes $\mathcal{C}^0 \cup \mathcal{C}^1 \dots \cup \mathcal{C}^t$ in the t -th session. Specifically, sufficient data can be provided in base session (also called initial session), while the model only has access to a few samples in all following sessions (also called incremental sessions). In particular, the incremental data are always organized as N -way K -shot format, *i.e.*, there are N classes, and each class contains K training images.

Several provident works [40, 51] recommend an incremental-frozen framework that leverages sufficient data in base session to train the model by optimizing the per-sample loss:


$$\mathcal{L}_{cls}(\phi; \mathbf{x}, y) = \mathcal{L}_{ce}(\phi(\mathbf{x}), y), \quad (1)$$

where $\mathcal{L}_{ce}(\cdot, \cdot)$ is cross-entropy loss (CE), $\phi(\mathbf{x})$ is the model which can be decomposed into feature extractor and classifier: $\phi(\mathbf{x}) = W^T f(\mathbf{x})$, where $\phi(\mathbf{x}) \in \mathbb{R}^{|\mathcal{C}^0| \times 1}$, $W \in \mathbb{R}^{d \times |\mathcal{C}^0|}$ and $f(\mathbf{x}) \in \mathbb{R}^{d \times 1}$.

After the base training stage, the backbone will be completely fixed, and the classifier is extended in each incremental session: $W = \{\mathbf{w}_1^0, \mathbf{w}_2^0, \dots, \mathbf{w}_{|\mathcal{C}^0|}^0\} \cup \dots \cup$



(a) Cumulative distribution function of inter-class distance. (b) Cumulative distribution function of intra-class distance.



(c) Base, novel and mutual class separation degree. (d) Base, novel and total classification accuracy.

Figure 2. The relationship between base class separation and novel class generalization. (a) Inter-class distance is defined in Eq. (5). A curve lies more right on the figure denotes greater inter-class distance. (b) Intra-class distance is defined in Eq. (6). A curve lies more left on the figure denotes smaller intra-class distance. (c) Class separation degree (R^2) is defined in Eq. (7).

$\{\mathbf{w}_1^t, \dots, \mathbf{w}_{|\mathcal{C}^t|}^t\}$. The classifier weight is parameterized by the average embeddings of each class (*i.e.*, prototype), represented as $\mathbf{w}_c^t = \frac{1}{n_c^t} \sum_{i=1}^{n_c^t} f(\mathbf{x}_{c,i})$.

In each session, the inference is conducted by the nearest class mean (NCM) algorithm [32] to evaluate accuracy of all encountered classes. Specifically, we measure the distance between $f(\mathbf{x})$ and all prototypes in classifiers:

$$c_{\mathbf{x}} = \arg \max_{c,t} \text{sim}(f(\mathbf{x}), \mathbf{w}_c^t), \quad (2)$$

where $\text{sim}(\mathbf{x}, \mathbf{y}) = \mathbf{x}^T \mathbf{y} / (\|\mathbf{x}\| \|\mathbf{y}\|)$ is cosine similarity between two vectors.

3.2. Base Class Separation Facilitates Novel Class Generalization

The incremental-frozen baseline outperforms a lot of incremental trainable methods [40], demonstrating that preserving old knowledge may matter more than accommodating new knowledge in FSCIL. The guarantee of base performance by the freezing strategy motivates our consideration of what our model should learn in the base training stage for future few-shot generalization.

An intuitive hypothesis is that, *base class separation facilitates novel class generalization*. If all prototypes are far away from each other, and the features of the same class

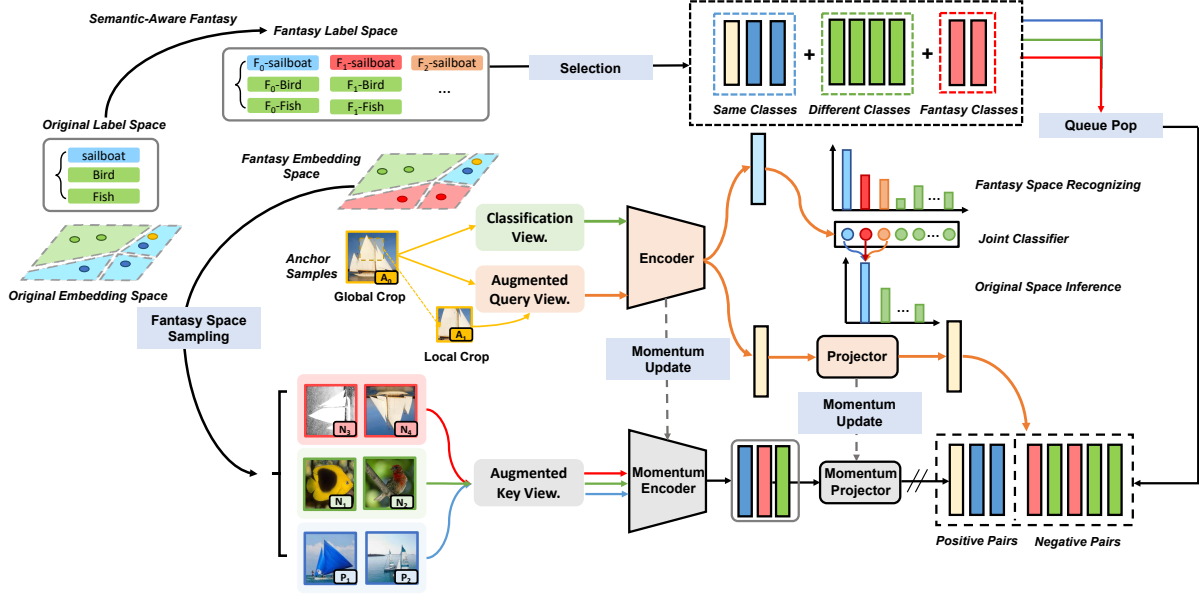


Figure 3. Overall pipeline of our Semantic-Aware Virtual Contrastive (SAVC) framework, which is composed of a classification branch and a contrastive learning branch. The backbone is shared for classification and contrastive query encoding, while the key backbone is updated with a large momentum (‘/’ means stop gradient). Both classification and contrastive loss are optimized in fantasy space, while the inference results of different transformations are aggregated and then projected into the original label space.

concentrate around their prototypes, there would be enough place for future class incorporation. Please note that the novel class generalization contains two meanings: without overlapping among novel class sets and with base class embeddings. Previous work demonstrated CE always cannot generalize well due to the poor class margins [13, 28] and may not be the optimal strategy for FSCIL, which inspires our attempt of supervised contrastive learning (SCL) [21], whose clustering effect has been demonstrated [53].

We adopt MoCo [7, 16] framework to implement SCL in our method, for the sake of lower memory budget and storing slowly evolving features of recent batches. For an given image-label pair (x, y) , we generate a query view $x_q = \text{Aug}_q(x)$ and a key view $x_k = \text{Aug}_k(x)$ by random augmentations. Then the two augmented images are fed into the query network g_q and key network g_k , getting two L_2 -normalized representations q and k , where $g = h \circ f$ is composed of a feature extractor f and a projector h . The two networks share the same architectures, while the key network updates itself with a large momentum. We maintain a feature queue and a label queue with the same length, storing the most current key embeddings and corresponding labels. To this end, given an image x , we have the following per-sample supervised contrastive loss:

$$\mathcal{L}_{cont}(g; \mathbf{x}, \tau, A) = -\frac{1}{|P(\mathbf{x})|} \sum_{\mathbf{k}_+ \in P(\mathbf{x})} \log \frac{\exp(\mathbf{q}^T \mathbf{k}_+ / \tau)}{\sum_{\mathbf{k}' \in A(\mathbf{x})} \exp(\mathbf{q}^T \mathbf{k}' / \tau)}, \quad (3)$$

where $A(\mathbf{x}) = \mathbf{k} \cup Q$ is the concatenation of the key embedding \mathbf{k} of \mathbf{x} and feature queue Q , $P(\mathbf{x})$ denotes positive set, *i.e.*, those in $A(\mathbf{x})$ belonging to the same class with \mathbf{x} .

In our practice, the feature extractor for query encoding and classification is shared. Put the training of the classification branch and contrastive learning branch together, we form a unified end-to-end framework, and the overall training objective can be concluded as:

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{cont}. \quad (4)$$

We would like to refer the model trained with Eq. (1) as CE, and the model trained with Eq. (4) as CE+SCL.

We then introduce and define some metrics to measure class separation and generalization, including:

1) Inter-class distance. Given two prototypes p_m and p_k of class m and k , we compute their distance and get $|\mathcal{C}^0| (|\mathcal{C}^0| - 1) / 2$ inter-class distances in base session.

$$d_{inter}^{m,k} = 1 - \text{sim}(p_m, p_k). \quad (5)$$

2) Intra-class distance. Given a class k , we compute the average Euclidean distance between all test samples belonging to class k and prototype p_k to measure the discrepancy within this class. For $|\mathcal{C}^0|$ classes in base session, we have:

$$d_{intra}^k = 1 - \frac{1}{n_k} \sum_{i=1}^{n_k} \text{sim}(f(\mathbf{x}_{k,i}), p_k). \quad (6)$$

3) Class Separation (R^2). On the one hand, we refer to Eq. (11) in [22] to measure the base ($K = |\mathcal{C}^0|$) and novel

($K = \sum_{t=1}^T |\mathcal{C}^t|$) class separation degree respectively:

$$\begin{aligned}
R^2 &= 1 - \bar{d}_{within} / \bar{d}_{total}, \\
\bar{d}_{within} &= \sum_{k=1}^K \sum_{i=1}^{n_k} \sum_{j=1}^{n_k} \frac{1 - \text{sim}(f(\mathbf{x}_{k,i}), f(\mathbf{x}_{k,j}))}{Kn_k^2}, \\
\bar{d}_{total} &= \sum_{k=1}^K \sum_{m=1}^K \sum_{i=1}^{n_k} \sum_{j=1}^{n_m} \frac{1 - \text{sim}(f(\mathbf{x}_{k,i}), f(\mathbf{x}_{m,j}))}{K^2 n_k n_j}.
\end{aligned} \tag{7}$$

On the other hand, we propose a novel metric to evaluate the mutual separation degree between novel and base classes, which is referred to as ‘mutual class separation (R_{mul}^2)’. We only need to let $K = \sum_{t=0}^T |\mathcal{C}^t|$ in \bar{d}_{within} , and $K = |\mathcal{C}^0|$ and $\sum_{t=1}^T |\mathcal{C}^t|$ respectively in \bar{d}_{total} , to compute R_{mul}^2 .

Greater d_{inter} and R_{base}^2 , smaller d_{intra} represents more compact the base embedding space is. Under the incremental-frozen framework, we compare the result of CE and CE+SCL in Fig. 2 on CIFAR100 benchmark, while d_{inter} and d_{intra} are shown in the cumulative distribution function (CDF) style. We can observe that although CE+SCL gets greater class separation degree and classification accuracy than CE on all classes, it implicitly decreases inter-class distance while decreasing intra-class distance (the CE+SCL curve lies on the left of the CE curve both in Fig. 2a and Fig. 2b). Especially, SCL only provides minor improvement of the mutual class separation R_{mul}^2 , which implies that SCL cannot separate the novel representations from the old distributions (*i.e.*, poor novel class generalization), and is also far from perfect for our objective (see Fig. 2c and Fig. 2d).

3.3. Semantic-Aware Contrastive Class Fantasy

The unsatisfactory performance of SCL, especially the inappreciable improvements of mutual class separation may convey that the semantic information base session provides is not diversified enough, which makes the learned decision boundary very sharp and close to data, resulting in poor base class separation and novel class generalization. Although the rests contained in new classes are invisible in base session, a provident model may be self-sufficient: imagining what would happen in the future and making adequate preparations for it. In other words, our model should use its fantasy to generate many virtual classes with various semantics and fill the unallocated embedding space, all by itself. These virtual classes act as not only class placeholders which preserve enough space for future updates, but also semantic knowledge providers which encourage extensive learning of different semantics for better generalization.

Semantic-aware class fantasy. In our semantic-aware contrastive class fantasy (SAVC) framework, we take inspiration from [24] and generate virtual classes by applying pre-defined transformations in self-supervised learning

(SSL) [11, 15, 35, 52]. In detail, we first define a discrete transformation (*i.e.*, fantasy) set \mathcal{F} and supposing there are M elements in it, *i.e.*, the number of transformations. Then, we can generate M transformed samples for an image-label pair (\mathbf{x}, y) , denoting $\mathcal{F}(\mathbf{x}, y) = \{(\mathbf{x}_m, y_m)\}_{m=1}^M$, where $y_m = y \times M + m$. The subscript of (\mathbf{x}_m, y_m) denotes m -th transformation is applied to (\mathbf{x}, y) , while (\mathbf{x}_1, y_1) represents the original image-label pair.

In this way, we enlarge our label space by M times, and each virtual class can be regarded as a ‘fine-grained’ class derived from the original one. The number and type of the elements in \mathcal{F} convey not only how compact we expect the embedding space to be, but also how fine the semantic grains we want to acquire. Hence we conduct the classification and contrasting training process in finer gains:

$$\begin{aligned}
\mathcal{L}_{cls}(\phi; \mathbf{x}, y, \mathcal{F}) &= \frac{1}{M} \sum_{i=1}^M \mathcal{L}_{ce}(\phi(\mathbf{x}_m), y_m), \\
\mathcal{L}_{cont}(g; \mathbf{x}, \tau, A, \mathcal{F}) &= \\
&= \frac{1}{M} \sum_{m=1}^M \frac{1}{|P(\mathbf{x}_m)|} \sum_{\mathbf{k}_+ \in P(\mathbf{x}_m)} \log \frac{\exp(\mathbf{q}^\top \mathbf{k}_+ / \tau)}{\sum_{\mathbf{k}' \in A(\mathbf{x}_m)} \exp(\mathbf{q}^\top \mathbf{k}' / \tau)}.
\end{aligned} \tag{8}$$

Multi-semantic aggregated inference. With the increase of class, we can acquire various derived class information in finer gains, which enables extending of classifiers with the prototypes of original and pseudo classes: $W = \{\mathbf{w}_{11}^0, \mathbf{w}_{12}^0, \dots, \mathbf{w}_{1M}^0, \dots, \mathbf{w}_{|\mathcal{C}^0|M}^0\} \cup \dots \cup \{\mathbf{w}_{11}^t, \dots, \mathbf{w}_{|\mathcal{C}^t|M}^t\}$. Prototypes with the same second subscript can naturally form a conditional prototype subset, which denotes a group of clustering centers aware of the same type of fine-grained semantics from the original classes: $W_m = \{\mathbf{w}_{1m}^0, \mathbf{w}_{2m}^0, \dots, \mathbf{w}_{|\mathcal{C}^0|m}^0\} \cup \dots \cup \{\mathbf{w}_{1m}^t, \dots, \mathbf{w}_{|\mathcal{C}^t|m}^t\}$. In contrast to the traditional single inference in Eq. (2), we first specify the test sample \mathbf{x} into various fine-grained semantics: $\mathcal{F}(\mathbf{x}) = \{\mathbf{x}_m\}_{m=1}^M$. Then we aggregate the inference results from different conditional prototype subsets, which improves accuracy significantly:

$$c_{\mathbf{x}} = \arg \max_{c,t} \sum_{m=1}^M \text{sim}(f(\mathbf{x}_m), \mathbf{w}_{cm}^t). \tag{9}$$

To this end, We compare our SAVC with CE and CE+SCL to verify its effectiveness. As we expected, SAVC effectively increases inter-class distance and decreases intra-class distance in base session, resulting in the best base class separation and novel class generalization (see Fig. 2).

3.4. Discussion

Our training pipeline is depicted in Fig. 3. Besides the training regime demonstrated in Sec. 3.3, the proposed SAVC also benefits from the following aspects:

Table 1. Comparison with SOTA methods on CUB200 dataset for FSCIL. *: Performances reported by [44]. Δ_{last} : Relative improvements of the last session compared to the Finetune baseline. More results on other datasets can be found in the appendix.

Method	Acc. in each session (%) \uparrow											Δ_{last}
	0	1	2	3	4	5	6	7	8	9	10	
Finetune* [44]	68.68	43.70	25.05	17.72	18.08	16.95	15.10	10.06	8.93	8.93	8.47	—
iCaRL* [37]	68.68	52.65	48.61	44.16	36.62	29.52	27.83	26.26	24.01	23.89	21.16	+12.69
EEIL* [4]	68.68	53.63	47.91	44.20	36.30	27.46	25.93	24.70	23.95	24.13	22.11	+13.64
TOPIC* [44]	68.68	62.49	54.81	49.99	45.25	41.40	38.35	35.36	32.22	28.31	26.28	+17.81
Rebalancing* [19]	68.68	57.12	44.21	28.78	26.71	25.66	24.62	21.52	20.12	20.06	19.87	+11.40
SPPR [58]	68.68	61.85	57.43	52.68	50.19	46.88	44.65	43.07	40.17	39.63	37.33	+28.86
F2M [40]	81.07	78.16	75.57	72.89	70.86	68.17	67.01	65.26	63.36	61.76	60.26	+51.79
CEC [51]	75.85	71.94	68.50	63.50	62.43	58.27	57.73	55.81	54.83	53.52	52.28	+43.81
FACT [55]	75.90	73.23	70.84	66.13	65.56	62.15	61.74	59.83	58.41	57.89	56.94	+48.47
SAVC (Ours)	81.85	77.92	74.95	70.21	69.96	67.02	66.16	65.30	63.84	63.15	62.50	+54.03




Figure 4. Comparison with SOTA methods on CIFAR100, *miniImageNet*, and CUB200 benchmarks.

Table 2. Statistics of benchmarks datasets. C^{base} : number of classes in base session. C^{inc} : total number of classes in incremental sessions. #Inc.: number of incremental sessions. Shots: training shots for incremental sessions.

Dataset	C^{base}	C^{inc}	#Inc.	Shots	Resolution
CIFAR100 [23]	60	40	8	5	32×32
<i>miniImageNet</i> [39]	60	40	8	5	84×84
CUB200 [48]	100	100	10	5	224×224

Multi-Crop strategy. Motivated by [3, 47], we add N constraint small crops to capture more local information, which is enforced to overlap with the global crops over a pre-defined threshold ensuring information sharing. These N smaller crops are only served as query views and not fed into the key encoder. In practice, we find optimizing the global and local contrastive loss respectively is helpful, and the overall training objective can be rewritten as:

$$\mathcal{L} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{cont.global} + \beta \mathcal{L}_{cont.local}. \quad (10)$$

Incremental finetuning. Although there are only a few novel samples available in incremental sessions, it is essential to finetune the pretrained model by Eq. (10) in each incremental session. Fortunately, we can contrast novel query samples with these slowly evolving key embeddings of base session from the feature queue, which helps push the novel

samples away from the distributions of the old classes and benefits novel class generalization.

4. Experiments

4.1. Experimental Setup

Datasets. Following the benchmark setting [51], we conduct experiments on three datasets, *i.e.*, CIFAR100 [23], *miniImageNet* [39] and CUB200 [48], to show the performance of our method. Their statistic characteristics are listed in Tab. 2.

Implementation details. Following [51], we adopt ResNet18 [17] backbones for *miniImageNet* and CUB200, and ResNet20 [17] for experiments on CIFAR100. We use SGD with 0.9 momentum to optimize the model. The initial learning rate is 0.1 for CIFAR100 and *miniImageNet*, and 0.002 for CUB200 in base session. For each incremental learning stage, we only finetune the model for 10 epochs in case of overfitting. We only consider 3 transformation sets \mathcal{F} , including 2-fold rotations, 4-fold rotations and 12-augmentations, and their concrete meanings are listed in Tab. 4. We use 12 augmentations for *miniImageNet*, and 2-fold rotation for CUB200 and CIFAR100, to fancy virtual classes. We set α as 0.2 and β as 0.8 for all datasets to control the impact of global and local contrastive loss.

Table 3. Ablation studies on *miniImageNet* benchmark. **SCL**, **F**, **MC**, **FT** denote supervised contrastive learning, virtual classes fantasy, multi-crop, incremental finetuning, respectively. Δ_{last} : Relative improvements of the last sessions compared to the CE baseline.

SCL	F	MC	FT	Acc. in each session (%) \uparrow								Δ_{last}	
				0	1	2	3	4	5	6	7		8
				70.27	65.17	61.13	57.80	54.80	51.84	49.14	47.03	45.54	—
✓				72.77	67.42	63.26	59.56	56.51	53.62	50.87	48.91	47.10	+1.56
✓	✓			80.72	75.82	71.46	67.65	64.63	61.31	58.27	56.20	54.52	+8.98
✓	✓	✓		81.12	76.02	71.81	68.41	65.46	62.11	59.33	57.42	55.80	+10.26
✓	✓	✓	✓	81.12	76.14	72.43	68.92	66.48	62.95	59.92	58.39	57.11	+11.57

Table 4. The concrete meaning of different types of transformation set \mathcal{F} . ‘ALL’ denotes using all 0° , 90° , 180° , 270° degrees.

\mathcal{F}	Rotation	Color permutation	M
2-fold rotations	0° , 180°	RGB	2
4-fold rotations	ALL	RGB	4
12-augmentations	ALL	RGB, GBR, BRG	12

4.2. Comparison with State of The Arts

We first compare our proposed SAVC with prevailing methods on three FSCIL benchmarks. These methods includes: classical CIL methods, *i.e.*, iCaRL [37], EEIL [4], Rebalancing [19], incremental-trainable FSCIL methods, *i.e.*, TOPIC [44], and incremental-frozen FSCIL methods, *i.e.*, SPPR [58], F2M [40], CEC [51], and FACT [55]. We also show a naïve baseline that directly finetunes the model with limited data as ‘finetune’. We report the performance curve in Fig. 4 and the detailed values on CUB200 in Tab. 1.

It is observed that ‘finetune’ leaves out preserving old knowledge and gets inferior performance. Classical CIL methods only consider extending the model with sufficient instances and are not suitable for FSCIL. Incremental-trainable FSCIL methods suffer from overfitting phenomena and perform poorly in incremental sessions. Incremental-frozen FSCIL methods can continually adapt to new classes without any training, which achieve better performance. Under the incremental-frozen framework, our SAVC effectively boosts both base class separation and novel class generalization by recognizing and contrasting semantic-aware virtual classes, consistently achieving the highest average accuracy over all sessions on three datasets. *e.g.*, our approach gets the final accuracies of 53.12%, 62.50% and 57.11% for CIFAR100, CUB200 and *miniImageNet*, outperforming the current SOTA method FACT by 1.02%, 5.56% and 6.62%, respectively.

4.3. Ablation study

We conduct ablation studies to prove the importance of our proposed components. Built on the incremental-frozen framework, we adopt only CE loss as the baseline, while adding SCL (SCL) and observe the influence this attempt




Figure 5. Comparison of the confusion matrices of CE and SAVC methods on CUB200 dataset.

brings (*i.e.*, CE and CE+SCL defined in Sec. 3.2). Then we fantasy (F) contrastive virtual class in CE+SCL and verify its effectiveness. In addition, we consider the influence of the two components mentioned in Sec. 3.4, *i.e.*, multi-crop strategy (MC) and finetune (FT) the pre-trained model. We report the results for *miniImageNet* dataset in Tab. 3.

Compared with the baseline, SCL shows limited influences on improving mutual class separation and lightly improves the performance by 1.56% in the last session. The fantasy of virtual classes (F) plays a vital part in preventing overlap among old and new distributions, which brings 8.98% performance improvement. When attending more local information (MC), our model works better, *i.e.*, 10.26% gains. Lastly, it is essential to fine-tune the model in incremental session (T), which achieves 11.57% accuracy gain.

4.4. Visualization

Confusion Matrix. We visualize the confusion matrix generated by CE and our SAVC on CUB200 dataset after the training of the last session in Fig. 5. Bright diagonal and dim background indicates higher classification accuracy. It is observed that CE draws a relatively distinct diagonal, but it has poor performance on new classes. In contrast, our method shows better performance on both base and new classes, which proves that our method adapts to novel classes without confusing the previous decision boundaries.

Visualization of class separation. We visualize the embedding space on CIFAR100 dataset with t-SNE [46] in Fig. 7. We randomly choose 6 classes from base classes and 3 classes from novel classes respectively. We can ob-




Figure 6. Analysis of detailed hyper configurations of our method: fantasy methods, trainable layers, and hyper-parameters.




Figure 7. The t-SNE [46] visualization on CIFAR100 dataset of the embeddings learned by various methods. Class 0-5 represent the base classes while class 6-8 represent the novel classes. Our SAVC gets the best class separation degree.

serve that CE has limited clustering effects. Naïve SCL slightly improves the class separation degree, but implicitly decreases inter-class distance and remains the mutual class confusion problems unresolved. With the help of virtual classes, Our proposed SAVC can clearly separate base classes and incorporate novel classes with the least overlap.

4.5. Further Analysis

Analysis of fantasy method. We consider 3 transformation sets \mathcal{F} and show their impacts on our method in Fig. 6a. We observe that *miniImageNet* benefits from complex transformation more, while it would hurt the class separation on CUB200 and CIFAR100. We suppose that CUB200 is already a fine-grained dataset, and the learning process would be hindered by overly fine grains. As for CIFAR100, the image resolutions are too small, and it may be hard to make a distinction between the original and virtual classes. We put discussions on more types of fantasy methods in the appendix.

Analysis of trainable layers in incremental sessions. We update the trainable layers in incremental sessions on *miniImageNet* benchmark. The results are reported in Fig. 6b. The gray horizontal line indicates without incremental finetuning. It is observed that finetuning fewer and deeper layers gets higher accuracy, while finetuning too

many layers harms the former knowledge and shows poorer performance. We infer that the low-level characteristics extracted by shallow layers are shared between classes, so we only need to update high layers to adapt to new classes.

Analysis of hyper-parameter. There are two main hyper-parameters in our loss, *i.e.*, coefficient α and β which control the impact of global and local supervised contrastive loss in Eq. (10). We report the accuracy in the last session on *miniImageNet* dataset varying α and β in Fig. 6c. Our method is robust to the choices of hyper-parameters, and it achieves the best performance when $\alpha = 0.2$ and $\beta = 0.8$. For convenience, we fix the value of α and β on all datasets.

5. Conclusion

In this paper, we solve the FSCIL problems by Semantic-Aware Virtual Contrastive framework (SAVC). Based on the mainstream incremental-frozen framework, both CE and SCL fail to separate base classes and new classes well. We propose to introduce more semantic knowledge by imagining virtual classes. These virtual classes not only hold enough feature space for future updates, but enable a multi-semantic aggregated inference effect and significantly improve classification accuracy. By learning to recognize and contrast in the fantasy space, our method can maintain the separation of all classes and achieve SOTA performance on the three popular FSCIL benchmarks.

Acknowledgments

This work is supported by grants from the Key-Area Research and Development Program of Guangdong Province (2020B0101380001), the National Key R&D Program of China (2022ZD0118101), the National Natural Science Foundation of China (61825101, 62202010, 62202014), Shenzhen Basic Research Program (JCYJ20220813151736001), China Postdoctoral Science Foundation (2022M710212), and Huawei Technologies Co., Ltd. The computing resources of Pengcheng Cloud-brain are used in this research.

References

- [1] Idan Achituve, Aviv Navon, Yochai Yemini, Gal Chechik, and Ethan Fetaya. Gp-tree: A gaussian process classifier for few-shot incremental learning. In *ICML*, pages 54–65. PMLR, 2021. 2
- [2] Touqeer Ahmad, Akshay Raj Dhamija, Steve Cruz, Ryan Rabinowitz, Chunchun Li, Mohsen Jafarzadeh, and Terrance E Boult. Few-shot class incremental learning leveraging self-supervised features. In *CVPRW*, pages 3900–3910, 2022. 2
- [3] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NeurIPS*, pages 9912–9924, 2020. 3, 6
- [4] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, pages 233–248, 2018. 6, 7
- [5] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with agem. *arXiv preprint arXiv:1812.00420*, 2018. 2
- [6] Kuilin Chen and Chi-Guhn Lee. Incremental few-shot learning via vector quantization in deep embedded space. In *ICLR*, 2020. 1, 2
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3, 4
- [8] Ali Cheraghian, Shafin Rahman, Pengfei Fang, Soumava Kumar Roy, Lars Petersson, and Mehrtash Harandi. Semantic-aware knowledge distillation for few-shot class-incremental learning. In *CVPR*, pages 2534–2543, 2021. 1, 2, 3
- [9] Ali Cheraghian, Shafin Rahman, Sameera Ramasinghe, Pengfei Fang, Christian Simon, Lars Petersson, and Mehrtash Harandi. Synthesized feature based few-shot class-incremental learning on a mixture of subspaces. In *ICCV*, pages 8661–8670, 2021. 2
- [10] 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 *CVPR*, pages 14166–14175, 2022. 2
- [11] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *ICCV*, pages 1422–1430, 2015. 2, 3, 5
- [12] Songlin Dong, Xiaopeng Hong, Xiaoyu Tao, Xinyuan Chang, Xing Wei, and Yihong Gong. Few-shot class-incremental learning via relation knowledge distillation. In *AAAI*, volume 35, pages 1255–1263, 2021. 1, 2, 3
- [13] Gamaleldin Elsayed, Dilip Krishnan, Hossein Mobahi, Kevin Regan, and Samy Bengio. Large margin deep networks for classification. In *NeurIPS*, 2018. 4
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, pages 1126–1135. PMLR, 2017. 2
- [15] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. *arXiv preprint arXiv:1803.07728*, 2018. 2, 3, 5
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9729–9738, 2020. 3, 4
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [18] Michael Hersche, Geethan Karunaratne, Giovanni Cherubini, Luca Benini, Abu Sebastian, and Abbas Rahimi. Constrained few-shot class-incremental learning. In *CVPR*, pages 9057–9067, 2022. 2
- [19] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *CVPR*, pages 831–839, 2019. 1, 6, 7
- [20] Zhong Ji, Zhishen Hou, Xiyao Liu, Yanwei Pang, and Xuelong Li. Memorizing complementation network for few-shot class-incremental learning. *TIP*, 2023. 3
- [21] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, pages 18661–18673, 2020. 2, 3, 4
- [22] Simon Kornblith, Ting Chen, Honglak Lee, and Mohammad Norouzi. Why do better loss functions lead to less transferable features? In *NeurIPS*, pages 28648–28662, 2021. 4
- [23] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [24] Hankook Lee, Sung Ju Hwang, and Jinwoo Shin. Self-supervised label augmentation via input transformations. In *ICML*, pages 5714–5724. PMLR, 2020. 3, 5
- [25] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 40(12):2935–2947, 2017. 1, 2
- [26] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few-shot learning. *arXiv preprint arXiv:1707.09835*, 2017. 2
- [27] 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 *ECCV*, pages 146–162. Springer, 2022. 2
- [28] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. Large-margin softmax loss for convolutional neural networks. *arXiv preprint arXiv:1612.02295*, 2016. 4
- [29] Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *ICPR*, pages 2262–2268. IEEE, 2018. 2
- [30] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost van de Weijer. Class-incremental learning: survey and performance evaluation on image classification. *arXiv preprint arXiv:2010.15277*, 2020. 1
- [31] Pratik Mazumder, Pravendra Singh, and Piyush Rai. Few-shot lifelong learning. In *AAAI*, volume 35, pages 2337–2345, 2021. 2
- [32] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *TPAMI*, 35(11):2624–2637, 2013. 3

- [33] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. [2](#)
- [34] Boris Oreshkin, Pau Rodríguez López, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018. [2](#)
- [35] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016. [2](#), [3](#), [5](#)
- [36] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. [2](#)
- [37] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, pages 2001–2010, 2017. [1](#), [2](#), [6](#), [7](#)
- [38] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. In *NeurIPS*, 2019. [2](#)
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115:211–252, 2015. [6](#)
- [40] Guangyuan Shi, Jiaxin Chen, Wenlong Zhang, Li-Ming Zhan, and Xiao-Ming Wu. Overcoming catastrophic forgetting in incremental few-shot learning by finding flat minima. In *NeurIPS*, pages 6747–6761, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [41] Yujun Shi, Kuangqi Zhou, Jian Liang, Zihang Jiang, Jiashi Feng, Philip HS Torr, Song Bai, and Vincent YF Tan. Mimicking the oracle: an initial phase decorrelation approach for class incremental learning. In *CVPR*, pages 16722–16731, 2022. [2](#)
- [42] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017. [2](#)
- [43] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, pages 1199–1208, 2018. [2](#)
- [44] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *CVPR*, pages 12183–12192, 2020. [1](#), [2](#), [3](#), [6](#), [7](#)
- [45] Eleni Triantafillou, Richard Zemel, and Raquel Urtasun. Few-shot learning through an information retrieval lens. In *NeurIPS*, 2017. [2](#)
- [46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. [7](#), [8](#)
- [47] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc V Gool. Revisiting contrastive methods for unsupervised learning of visual representations. In *NeurIPS*, pages 16238–16250, 2021. [6](#)
- [48] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [6](#)
- [49] Zifeng Wang, Zizhao Zhang, Chen-Yu Lee, Han Zhang, Ruoxi Sun, Xiaoqi Ren, Guolong Su, Vincent Perot, Jennifer Dy, and Tomas Pfister. Learning to prompt for continual learning. In *CVPR*, pages 139–149, 2022. [2](#)
- [50] Shipeng Yan, Jiangwei Xie, and Xuming He. Der: Dynamically expandable representation for class incremental learning. In *CVPR*, pages 3014–3023, 2021. [1](#), [2](#)
- [51] Chi Zhang, Nan Song, Guosheng Lin, Yun Zheng, Pan Pan, and Yinghui Xu. Few-shot incremental learning with continually evolved classifiers. In *CVPR*, pages 12455–12464, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [52] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, pages 649–666. Springer, 2016. [2](#), [3](#), [5](#)
- [53] Yifan Zhang, Bryan Hooi, Dapeng Hu, Jian Liang, and Jiashi Feng. Unleashing the power of contrastive self-supervised visual models via contrast-regularized fine-tuning. In *NeurIPS*, pages 29848–29860, 2021. [4](#)
- [54] Hanbin Zhao, Yongjian Fu, Mintong Kang, Qi Tian, Fei Wu, and Xi Li. Mgsvf: Multi-grained slow vs. fast framework for few-shot class-incremental learning. *TPAMI*, 2021. [2](#), [3](#)
- [55] 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 *CVPR*, pages 9046–9056, 2022. [1](#), [3](#), [6](#), [7](#)
- [56] 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. *TPAMI*, 2022. [3](#)
- [57] Fei Zhu, Zhen Cheng, Xu-Yao Zhang, and Cheng-lin Liu. Class-incremental learning via dual augmentation. In *NeurIPS*, pages 14306–14318, 2021. [2](#)
- [58] Kai Zhu, Yang Cao, Wei Zhai, Jie Cheng, and Zheng-Jun Zha. Self-promoted prototype refinement for few-shot class-incremental learning. In *CVPR*, pages 6801–6810, 2021. [2](#), [3](#), [6](#), [7](#)