

# Tree Structure-Aware Few-Shot Image Classification via Hierarchical Aggregation

Min Zhang<sup>124</sup> Siteng Huang<sup>24</sup> Wenbin Li<sup>3</sup> Donglin Wang<sup>24\*</sup>

<sup>1</sup>Zhejiang University <sup>2</sup>Westlake University

<sup>3</sup>State Key Laboratory for Novel Software Technology, Nanjing University

<sup>4</sup>Institute of Advanced Technology, Westlake Institute for Advanced Study  
liwenbin@nju.edu.cn, {zhangmin, huangsiteng, wangdonglin}@westlake.edu.cn

**Abstract.** In this paper, we mainly focus on the problem of how to learn additional feature representations for few-shot image classification through pretext tasks (*e.g.*, rotation or color permutation and so on). This additional knowledge generated by pretext tasks can further improve the performance of few-shot learning (FSL) as it differs from human-annotated supervision (*i.e.*, class labels of FSL tasks). To solve this problem, we present a plug-in *Hierarchical Tree Structure-aware (HTS)* method, which not only learns the relationship of FSL and pretext tasks, but more importantly, can adaptively select and aggregate feature representations generated by pretext tasks to maximize the performance of FSL tasks. A hierarchical tree constructing component and a gated selection aggregating component is introduced to construct the tree structure and find richer transferable knowledge that can rapidly adapt to novel classes with a few labeled images. Extensive experiments show that our HTS can significantly enhance multiple few-shot methods to achieve new state-of-the-art performance on four benchmark datasets. The code is available at: <https://github.com/remiMZ/HTS-ECCV22>.

**Keywords:** Hierarchical Tree Structure, Few-shot Learning, Pretext Tasks

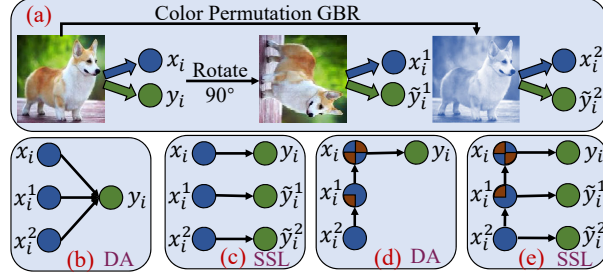
## 1 Introduction

Few-shot learning (FSL), especially few-shot image classification [17, 37, 2, 22, 44], has attracted a lot of machine learning community. FSL aims to learn transferable feature representations by training the model with a collection of FSL tasks on base (seen) classes and generalizing the representations to novel (unseen) classes by accessing an extremely few labeled images [1, 7, 38, 10, 30, 4]. However, due to the data scarcity, the learned supervised representations mainly focus on the differences between the base class while ignoring the valuable semantic features within images for novel classes, weakening the model’s generalization ability. Therefore, more feature representations should be extracted from the limited available images to improve the generalization ability of the FSL model.

One effective way of extracting more useful feature representations is to use pretext tasks, such as rotation with multiple angles or color permutation among

---

\* Corresponding author



**Fig. 1.** Differences in the learning process for few-shot image classification using pretext tasks between previous and our works. (a) shows the process of generating augmented images using FSL images. (b) and (c) show the learning process for previous works under the DA or SSL setting, which uses all images indiscriminately. (d) and (e) show the learning process for our work under the DA or SSL setting, which can exploit the hierarchical tree structure to adaptively select useful feature representations.

different channels [11,16,34,24,5]. Because these pretext tasks can generate additional augmented images and the semantic representations of these augmented images are a good supplementary to the normal human-annotated supervisions (*i.e.*, class labels of FSL images), which is beneficial to the few-shot learning model generalization on novel classes. The standard training process of using pretext tasks to assist in FSL can be roughly classified into two settings, *i.e.*, data augmentation (DA) and self-supervised learning (SSL), following existing works [11,34,26,24,37]. As shown in Fig. 1, (a) shows that pretext tasks are used to generate multiple augmented images ( $x_i^1, x_i^2$ ) and  $x_i$  is FSL images. (b) and (c) are the learning process of using pretext tasks to improve the FSL performance under the DA or SSL setting in previous works. However, in the DA setting, all augmented and raw images are placed in the same label space (*e.g.*,  $y_i = \tilde{y}_i^1 = \tilde{y}_i^2 = \text{dog}$ ) and the empirical risk minimization (ERM) (see Eq. (3)) is used to optimize the model, making the model use all images indiscriminately.

We find that when augmented images are generated by using inappropriate pretext tasks, this optimization method (*i.e.*, considering the information of all images on average) may destroy the performance of the FSL task (see Fig. 3). This is because the augmented images bring ambiguous semantic information (*e.g.*, rotation for symmetrical object) [9,6,25]. Although it can be solved using expert experience to select appropriate pretext tasks for different datasets, which is very labor consuming [26]. To this end, we believe that it is very important that the model can adaptively select augmented image features to improve the performance of the FSL task. In the SSL setting (Fig. 1 (c)), it retains FSL as the main task and uses pretext tasks as additional auxiliary tasks (*i.e.*, SSL tasks). From Fig. 3, we find that SSL using independent label spaces (*e.g.*,  $y_i = \text{dog}$ ,  $\tilde{y}_i^1 = 90^\circ$ ,  $\tilde{y}_i^2 = \text{GBR}$ ) to learn these tasks separately (see Eq. (4)) can alleviate the problem caused by DA training based on one single label space, but it is not enough to fully learn the knowledge hidden in these augmented images only through the sharing network. This is because there are similarities among augmented images generated under the same raw image and different pretext tasks, and the relationships among these augmented images should be modeled.

To effectively learn knowledge from augmented images, we propose a plug-in *Hierarchical Tree Structure-aware (HTS)* method for few-shot classification. The core of the proposed HTS method: (1) using a tree structure to model the relationships of raw and augmented images; (2) using a gated aggregator to adaptively select the feature representations to improve the performance of the FSL task. Next, we outline two key components of the proposed HTS method.

**Modeling the relationships via a hierarchical tree constructing component.** This component aims to construct a tree structure for each raw image, and thus we can use the edge of the tree to connect the feature information among different augmented images and use the level of the tree to learn the feature representations from different pretext tasks. In addition, when pretext tasks or augmented images change (*e.g.*, adding, deleting or modifying), it is very flexible for our HTS method to change the number of levels or nodes.

**Adaptively learning features via a gated selection aggregating component.** In this paper, we use *Tree-based Long Short Term Memory (TreeLSTM)* [36] as the gated aggregator following the reasons as below: (1) On the above trees, we find that the augmented images (*i.e.*, nodes) from different pretext tasks can be further formulated as a sequence with variable lengths from the bottom to the top level. (2) TreeLSTM generates a forgetting gate for each child node, which is used to filter the information of the corresponding child nodes (different colors are shown in Fig. 1 (d) and (e)). This indicates that the representations of the lower-level nodes can be sequentially aggregated and enhance the upper-level nodes’ outputs. Finally, these aggregated representations will be used in training and testing phases. The main contributions of HTS are:

1. We point out the limitations of using pretext tasks to help the few-shot model learn richer and transferable feature representations. To solve these limitations, we propose a hierarchical tree structure-aware method.
2. We propose a hierarchical tree constructing component to model the relationships of augmented and raw images and a gated selection aggregating component to adaptively learn and improve the performance of the FSL task.
3. Extensive experiments on four benchmark datasets demonstrate that the proposed HTS is significantly superior to the state-of-the-art FSL methods.

## 2 Related Work

### 2.1 Few-Shot Learning

The recent few-shot learning works are dominated by meta-learning based methods. They can be roughly categorized into two groups: (1) *Optimization-based methods* advocate learning a suitable initialization of model parameters from base classes and transferring these parameters to novel classes in a few gradient steps [31,45,3,10,27]. (2) *Metric-based methods* learn to exploit the feature similarities by embedding all images into a common metric space and using well-designed nearest neighbor classifiers [19,20,14,42]. In this paper, our HTS can equip with an arbitrary meta-learning based method and improve performance.

## 2.2 Pretext Tasks

Pretext tasks have succeeded to learn useful representations by focusing on richer semantic information of images to significantly improve the performance of image classification. In this paper, we are mainly concerned with the works of using pretext tasks to improve the performance of few-shot classification [11,12,43,16,5]. However, these works are often shallow, *e.g.*, the original FSL training pipeline is intact and an additional loss (self-supervised loss) on each image is introduced, leading to the learning process not being able to fully exploit the augmented image representations. Different from these works, we introduce a hierarchical tree structure (HTS) to learn the pretext tasks. Specifically, the relationships of each image are modeled to learn more semantic knowledge. Moreover, HTS can adaptively select augmented features to avoid the interference of ambiguous information. Our experimental results show that thanks to the reasonable learning of pretext tasks, our HTS method clearly outperforms these works (see Tab. 2).

## 3 Preliminaries

### 3.1 Problem Setting in Few-Shot Learning

We consider that meta-learning based methods are used to solve the few-shot classification problem, and thus follow the episodic (or task) training paradigm. In the meta-training phase, we randomly sample episodes from a base class set  $\mathcal{D}_b$  to imitate the meta-testing phase sampled episodes from a novel class set  $\mathcal{D}_n$ . Note that  $\mathcal{D}_b$  contains a large number of labeled images and classes but has a *disjoint* label space with  $\mathcal{D}_n$  (*i.e.*  $\mathcal{D}_b \cap \mathcal{D}_n = \emptyset$ ). Each  $n$ -way  $k$ -shot episode  $\mathcal{T}_e$  contains a support set  $\mathcal{S}_e$  and a query set  $\mathcal{Q}_e$ . Concretely, we first randomly sample a set of  $n$  classes  $\mathcal{C}_e$  from  $\mathcal{D}_b$ , and then generate  $\mathcal{S}_e = \{(x_i, y_i) | y_i \in \mathcal{C}_e, i = 1, \dots, n \times k\}$  and  $\mathcal{Q}_e = \{(x_i, y_i) | y_i \in \mathcal{C}_e, i = 1, \dots, n \times q\}$  by sampling  $k$  support and  $q$  query images from each class in  $\mathcal{C}_e$ , and  $\mathcal{S}_e \cap \mathcal{Q}_e = \emptyset$ . For simplicity, we denote  $l_k = n \times k$  and  $l_q = n \times q$ . In the meta-testing phase, the trained few-shot learning model is fine-tuned using the support set  $\mathcal{S}_e$  and is tested using the query set  $\mathcal{Q}_e$ , where the two sets are sampled from the novel class set  $\mathcal{D}_n$ .

### 3.2 Few-Shot Learning Classifier

We employ ProtoNet [33] as the few-shot learning (FSL) model for the main instantiation of our HTS framework due to its simplicity and popularity. However, we also show that any meta-learning based FSL method can be combined with our proposed HTS method (see results in Tab. 3). ProtoNet contains a feature encoder  $E_\phi$  with learnable parameters  $\phi$  (*e.g.*, CNN) and a simple non-parametric classifier. In each episode  $\mathcal{T}_e = \{\mathcal{S}_e, \mathcal{Q}_e\}$ , ProtoNet computes the mean feature embedding of the support set for each class  $c \in \mathcal{C}_e$  as the prototype  $\tilde{p}_c$ :

$$\tilde{p}_c = \frac{1}{k} \sum_{(x_i, y_i) \in \mathcal{S}_e} E_\phi(x_i) \cdot \mathbb{I}(y_i = c), \quad (1)$$

where  $\mathbb{I}$  is the indicator function with its output being 1 if the input is true or 0 otherwise. Once the class prototypes are obtained from the support set, ProtoNet computes the distances between the feature embedding of each query

set image and that of the corresponding prototypes. The final loss function over each episode using the empirical risk minimization (ERM) is defined as follows:

$$L_{FSL}(\mathcal{S}_e, \mathcal{Q}_e) = \frac{1}{|\mathcal{Q}_e|} \sum_{(x_i, y_i) \in \mathcal{Q}_e} -\log p_{y_i},$$

$$p_{y_i} = \frac{\exp(-d(E_\phi(x_i), \tilde{p}_{y_i}))}{\sum_{c \in \mathcal{C}_e} \exp(-d(E_\phi(x_i), \tilde{p}_c))}, \quad (2)$$

where  $d(\cdot, \cdot)$  denotes a distance function (*e.g.*, the squared euclidean distance for ProtoNet method following the original paper [33]).

## 4 Methodology

### 4.1 Pretext Tasks in FSL

Pretext tasks assisting in few-shot learning have two settings: data augmentation (DA) and self-supervised learning (SSL) (see the schematic in Fig. 1). We first define a set of pretext-task operators  $\mathcal{G} = \{g_j | j = 1, \dots, J\}$ , where  $g_j$  means the operator of using the  $j$ -th pretext task and  $J$  is the total number of pretext tasks. Moreover, we also use  $M_j$  to represent the number of augmented images generated by using the  $j$ -th pretext task for each raw image and the pseudo label set of this task is defined as  $\tilde{Y}^j = \{0, \dots, M_j - 1\}$ . For example, for a 2D-rotation operator, each raw image will be rotated with multiples of  $90^\circ$  angles (*e.g.*,  $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ), where the augmented images are  $M_{rotation} = 3$  and the pseudo label set is  $\tilde{Y}^{rotation} = \{0, 1, 2\}$ . Given a raw episode  $\mathcal{T}_e = \{\mathcal{S}_e, \mathcal{Q}_e\}$  as described in Sec. 3.1, we utilize these pretext-task operators from  $\mathcal{G}$  in turn to augment each image in  $\mathcal{T}_e$ . This results in a set of  $J$  augmented episodes are  $\mathcal{T}_{aug} = \{(x_i, y_i, \tilde{y}_i, j) | y_i \in \mathcal{C}_e, \tilde{y}_i \in \tilde{Y}^j, i = 1, \dots, M_j \times l_k, M_j \times (l_k + 1), \dots, M_j \times (l_k + l_q), j = 1, \dots, J\}$ , where the first images  $M_j \times l_k$  are from the augmented support set  $\mathcal{S}_e^j$  and the rest of images  $M_j \times l_q$  from the augmented query set  $\mathcal{Q}_e^j$ .

**Data Augmentation.** For DA setting, we use the combined episodes  $\mathcal{T} = \{\{\mathcal{S}_e^r, \mathcal{Q}_e^r\} | r = 0, \dots, J\}$ , where  $\{\mathcal{S}_e^0, \mathcal{Q}_e^0\}$  is the raw episode and  $\{\{\mathcal{S}_e^r, \mathcal{Q}_e^r\} | r = 1, \dots, J\}$  is the augmented episodes. In this paper, when  $r \geq 1$ , the value of  $r$  is equal to  $j$  unless otherwise stated. Each image  $(x_i, y_i)$  in  $\mathcal{T}$  takes the same class label  $y_i$  (from the human annotation) for supervised learning to improve the performance of the FSL. The objective is to minimize a cross-entropy loss:

$$L_{DA} = \frac{1}{J+1} \sum_{r=0}^J L_{FSL}(\mathcal{S}_e^r, \mathcal{Q}_e^r). \quad (3)$$

$L_{DA}$  uses the empirical risk minimization (ERM) algorithm based on the same label space (*e.g.*,  $y_i$ ) to learn raw and augmented feature representations. However, if the augmented images have ambiguous representations, this optimization method may interfere with the semantic learning of the FSL model.

**Self-Supervised Learning.** For SSL setting, each raw image  $(x_i, y_i)$  in  $\mathcal{T}_e$  uses a class label  $y_i$  for supervised learning, while each augmented image  $(x_i, \tilde{y}_i)$  in  $\mathcal{T}_{aug}$  carries a pseudo label  $\tilde{y}_i$  for self-supervised learning. A multi-task learning loss (FSL main task and SSL auxiliary task) is normally adopted as below:

$$L_{SSL} = L_{FSL}(\mathcal{S}_e, \mathcal{Q}_e) + \sum_{j=1}^J \beta_j L_j,$$

$$L_j = \frac{1}{E} \sum_{(x_i, \tilde{y}_i \in \mathcal{T}_e^j)} -\log \frac{\exp([\theta^j(E_\phi(x_i))]_{\tilde{y}_i})}{\sum_{\tilde{y}' \in \tilde{\mathcal{Y}}^j} \exp([\theta^j(E_\phi(x_i))]_{\tilde{y}'}), \quad (4)$$

where  $E = M_j \times (l_k + l_q)$ ,  $[\theta^j(E_\phi(x_i))]$  denotes the  $j$ -th pretext task scoring vector and  $[\cdot]_{\tilde{y}}$  means taking the  $\tilde{y}$ -th element.  $L_{SSL}$  learns the knowledge of the few-shot learning and multiple sets of pretext tasks in different label spaces, but only uses a shared feature encoder  $E_\phi$  to exchange these semantic information.

## 4.2 Pretext Tasks in HTS

In this paper, we propose a hierarchical tree structure-aware (HTS) method that uses a tree to model relationships and adaptively selects knowledge among different image features. There are two key components in HTS: hierarchical tree constructing component and gated selection aggregating component <sup>1</sup>.

### 4.2.1 Hierarchical Tree Constructing Component


Given augmented episodes  $\mathcal{T}_{aug} = \{\mathcal{T}_e^j = \{\mathcal{S}_e^j, \mathcal{Q}_e^j\} | j = 1, \dots, J\}$  as described in Sec. 4.1, each augmented episode in  $\mathcal{T}_{aug}$  corresponds to one specific pretext task of the same set of images from the raw episode  $\mathcal{T}_e$ . Therefore, we believe that these augmented images with different pretext tasks should be modeled to capture the correlations and further learn more semantic feature representations.

To this end, we construct a tree structure for each FSL image and its corresponding multiple sets of augmented images generated by using different pretext tasks in each episode. Specifically, (1) we extract the feature vectors of the raw and augmented episodes by using the shared feature encoder  $E_\phi$  and denote the feature set as  $\mathcal{T}_{emd}$ , where  $\mathcal{T}_{emd} = \{E_\phi(x_i) | (x_i, y_i, \tilde{y}_i, r) \in \mathcal{T}_e^r, r = 0, \dots, J, i = 1, \dots, M_j \times (l_k + l_q)\}$ . (2) The feature vectors of the raw episode  $\mathcal{T}_{emd}^0$  are taken as the root nodes and each raw image has its own tree structure. (3) The augmented feature vectors of these augmented episodes  $\mathcal{T}_{emd}^j$  with the  $j$ -th pretext task are put in the  $(j+1)$ -th level of the tree. (4) We take a raw image  $x_i$  (i.e., a tree structure) and its multiple sets of augmented images  $\{x_i^j\}$  as an example to indicate how to construct this tree, and repeat the process for other raw images. The form of the tree is  $\{E_\phi(x_i) \xrightarrow{g_1} E_\phi(x_i^1) \xrightarrow{g_2} \dots \xrightarrow{g_j} E_\phi(x_i^j) \dots \xrightarrow{g_J} E_\phi(x_i^J)\}$ , where the raw feature set  $E_\phi(x_i)$  is in the 1-st level (root nodes) and the augmented feature set is in the  $(j+1)$ -th level sharing the same pretext task  $g_j$ . For each episode, we construct  $(l_k + l_q)$  hierarchical tree structures, and every level has  $M_j$  child nodes with the  $j$ -th pretext task. In these tree structures, the edge information is used to model the relationships of different augmented or raw images. The level knowledge is used to learn the representations from different pretext tasks. In Sec. 4.2.2, we introduce how to better aggregate image features.

### 4.2.2 Gated Selection Aggregating Component

As mentioned above, we have constructed a tree structure for each raw image

<sup>1</sup> Note that our method mainly focuses on how to adaptively learn the knowledge of pretext tasks and improve the performance of few-shot image classification.



**Fig. 2.** The learning process of gated selection aggregating component. (a) shows that the aggregator to sequentially and hierarchically aggregate the information from bottom to top level. (b) details the internal aggregation of TreeLSTM (e.g., two levels). The subscript marked in this figure represents the number of child nodes, the superscript represents the number of levels and different colors are different LSTM cells.

$x_i$  in each randomly sampled episode. The following question is how to effectively use this tree structure for learning and inference. Firstly, our intuition is to preserve the tree structure information, because it models the relationships among images (e.g., FSL and augmentations). Secondly, we should selectively aggregate the features of all child nodes from the bottom to the top level, because the information aggregated process aims to maximize the performance of the parent node after aggregation. Thirdly, since the feature information in the hierarchical propagation can be regarded as sequential inputs divided by levels but meta-learning can not directly process the sequential data. Finally, we adopt the *tree-based Long Short Term Memory (TreeLSTM)* as our gated aggregator to encode the lower-level information into an upper-level output. In this way, we can mine much richer features from the tree structure (see Fig. 2). Next, we will detail the aggregation and propagation process of TreeLSTM on these trees.

For simplicity, we take a tree structure as an example to introduce how to selectively aggregate information by using TreeLSTM aggregator and repeat the process for other tree structures. The  $(J+1)$ -level tree structure are constructed based on raw image features  $E_\phi(x_i)$  and its multiple sets of augmented features  $\{E_\phi(x_i^j) | i = 1, \dots, M_j, j = 1, \dots, J\}$ . The TreeLSTM aggregates information step by step from bottom level (i.e.,  $(J+1)$ -th level) to top level (i.e., 1-st or root node level). We use  $\{h_i^r | r = 0, \dots, J\}$  to represent the aggregated node representations of each level except the bottom-level nodes in this tree. Because the bottom node has no child nodes, its aggregate information is itself  $\{E_\phi(x_i^J)\}$ . The aggregation process can be formalized as:  $\{h_i^0 \xleftarrow{agg} h_i^1 \xleftarrow{agg} \dots \xleftarrow{agg} h_i^r \dots \xleftarrow{agg} E_\phi(x_i^J)\}$ , where *agg* denotes the aggregation operation by using the TreeLSTM. The aggregated output  $h_i^r$  of each level is represented as:

$$h_i^r = \text{TreeLSTM}(s_i, \{h_m\}), m \in \mathcal{M}_i, \quad (5)$$

where  $s_i \in \{h_i^r\}$  is any node in the tree and  $\mathcal{M}_i$  is the set of child nodes of the  $i$ -th node in  $(r+1)$  level of the tree.

Since this form of child-sum in TreeLSTM conditions its components on the sum of child hidden states  $h_m$ , it is a permutation invariant function and is well-suited for trees whose children are unordered. But, we find that using the child-



mean replaces the child-sum in our implementation for better normalization, compared with the original paper [36]. The formulation of the TreeLSTM is:

$$\begin{aligned}
f_m &= \sigma(W_f s_i + U_f h_m + b_f), \\
h_{me} &= \frac{\sum_{m \in \mathcal{M}_i} h_m}{|\mathcal{M}_i|}, \quad u_i = \tanh(W_u s_i + U_u h_{me} + b_u), \\
o_i &= \sigma(W_o s_i + U_o h_{me} + b_o), \quad i_i = \sigma(W_i s_i + U_i h_{me} + b_i), \\
c_i &= i_i \odot u_i + \frac{\sum_{m \in \mathcal{M}_i} f_m \odot c_m}{|\mathcal{M}_i|}, \quad h_i^r = o_i \odot \tanh(c_i),
\end{aligned} \tag{6}$$

where  $\odot$  denotes the element-wise multiplication,  $\sigma$  is sigmoid function,  $W_a, U_a, b_a$ ,  $a \in \{i, o, f, u\}$  are trainable parameters,  $c_i, c_k$  are memory cells,  $h_i, h_m$  are hidden states and  $h_{me}$  denotes the child nodes mean in  $(r+1)$ -th level. From Eq. (6) and Fig. 2, the aggregator learns a forgetting gate for each child and uses it to adaptively select the child features that are beneficial to its parent node based on the trained parameters. Finally, the aggregated episodes are denoted as  $\mathcal{T}_{agg} = \{\{\mathcal{S}_{agg}^r, \mathcal{Q}_{agg}^r\} | r = 0, \dots, J\}$ , where  $\mathcal{S}_{agg}^r = \{h_i^r | i = 1, \dots, M_r * l_k\}$  and  $\mathcal{Q}_{agg}^r = \{h_i^r | i = 1, \dots, M_r * l_q\}$  with  $M_0 = 1, M_r = M_j$  when  $r \geq 1$ .

### 4.3 Meta-Training Phase with HTS

The meta-training models with HTS use previous settings (DA or SSL) to train the FSL model. Given the aggregated episodes  $\mathcal{T}_{agg} = \{\{\mathcal{S}_{agg}^r, \mathcal{Q}_{agg}^r\} | r = 0, \dots, J\}$ , where  $\{\mathcal{S}_{agg}^0, \mathcal{Q}_{agg}^0\}$  is the raw representations aggregated by using its all augmented images (*i.e.*, from  $(J+1)$ -th to 2-nd of the tree) and  $\{\mathcal{S}_{agg}^r, \mathcal{Q}_{agg}^r\}_{r=1}^J$  is the augmented representations with each pretext task aggregated by using other augmented images (*e.g.*, from  $(J+1)$ -th to  $r$ -th of the tree).

**Data Augmentation.** For data augmentation (DA) setting, we only use the aggregated root node (raw images)  $\{\mathcal{S}_{agg}^0, \mathcal{Q}_{agg}^0\}$  to train the FSL model. The cross-entropy loss can be defined as:

$$L_{DA}^{HTS} = L_{FSL}(\mathcal{S}_{agg}^0, \mathcal{Q}_{agg}^0). \tag{7}$$

$L_{HTS}^{DA}$  uses the aggregated representations with more knowledge to train the few-shot image classification model, different from Eq. (3).

**Self-Supervised Learning.** For SSL setting, the aggregated root nodes train the FSL main task and each aggregated augmented node  $(h_i^r, \tilde{y}_i)$  in  $\{\mathcal{T}_{agg}^r = \{\mathcal{S}_{agg}^r, \mathcal{Q}_{agg}^r\} | r = 1, \dots, J\}$  trains the SSL auxiliary task. With the pseudo label  $\tilde{y}_i$  and every level aggregated features in the tree, the multi-task learning loss is:

$$\begin{aligned}
L_{SSL}^{HTS} &= L_{FSL}(\mathcal{S}_{agg}^0, \mathcal{Q}_{agg}^0) + \sum_{r=1}^J \beta_r L_r, \\
L_r &= \frac{1}{E} \sum_{(h_i^r, \tilde{y}_i) \in \mathcal{T}_{agg}^r} -\log \frac{\exp([\theta^r(h_i^r)]_{\tilde{y}_i})}{\sum_{\tilde{y}' \in \tilde{\mathcal{Y}}^J} \exp([\theta^r(h_i^r)]_{\tilde{y}'}),
\end{aligned} \tag{8}$$

where  $E$  and  $[\theta^r(\cdot)]_{\tilde{y}}$  have the same meanings as Eq. (4). and when  $r \geq 1$ ,  $\theta_r = \theta_j$ . For easy reproduction, we summarize the overall algorithm for FSL with HTS in Algorithm 1 in Appendix A.1. Notably, our HTS is a lightweight



method, consisting of a parameter-free tree constructing component and a simple gated aggregator component. Therefore, as a plug-in method, HTS will not introduce too much additional computing overhead<sup>2</sup>. Once trained, with the learned network parameters, we can perform the testing over the test episodes.

#### 4.4 Meta-Testing Phase with HTS

During the meta-testing stage, we find that using pretext tasks for both the query set and support set, and then aggregating features based on the trained aggregator can further bring gratifying performance with only a few raw images (see Fig. 4). Therefore, the predicted label for  $x_i \in \mathcal{Q}_e$  can be computed with Eq. (1) and aggregated raw features  $h_i^0$  as:

$$y_i^{pred} = \underset{y \in \mathcal{C}_e}{argmax} \frac{\exp(-d(h_i^0, \tilde{p}_{y_i}))}{\sum_{c \in \mathcal{C}_e} \exp(-d(h_i^0, \tilde{p}_c))}. \quad (9)$$

#### 4.5 Connections to Prior Works

**Hierarchical few-shot learning.** Recently, some works have been proposed to solve the problem of few-shot learning using the hierarchical structure. One representative work aims to improve the effectiveness of meta-learning by hierarchically clustering different tasks based on the level of model parameter [40], while our method is based on the feature level with less running time and computation. Another work [23] learns the relationship between fine-grained and coarse-grained images through hierarchical modeling, but it requires the data itself to be hierarchical, while our method can adapt to any dataset.

**Long short term memory few-shot learning.** Some works [38, 41, 21] use chain-based long short-term memory (ChainLSTM) to learn the relationship of images. However, our work uses tree-based long short-term memory (TreeLSTM) to learn structured features and meanwhile preserve tree-structured information.


## 5 Experimental Results

### 5.1 Experimental Setup

**Research Questions.** Research questions guiding the remainder of the paper are as follows: **RQ1.** What is the real performance of the heuristic combination of FSL main task and pretext auxiliary tasks? **RQ2.** How effective is the proposed HTS framework for the FSL image classification task in both single-domain and cross-domain settings? **RQ3.** Could the proposed HTS framework adaptively select augmented features for better improving the performance of the FSL main task? **RQ4.** How does the proposed HTS method work (ablation study)?

**Benchmark Datasets.** All experiments are conducted on four FSL benchmark datasets, *i.e.*, *miniImageNet* [38], *tieredImageNet* [29], *CUB-200-2011* [39] and

<sup>2</sup> During training, for a 5-way 1-/5-shot setting, one episode time is 0.45/0.54s (0.39/0.50s for baseline) with 75 query images over 500 randomly sampled episodes.



**Fig. 3.** The results (5-shot) for motivation on *miniImageNet* (MINI), CUB-200-2011 (CUB), CIFAR-FS (CIFAR) and *tieredImageNet* (TIERED) with two pretext tasks.

CIFAR-FS [3]. The *miniImageNet* and *tieredImageNet* are the subsets of the ILSVRC-12 dataset [8]. CUB-200-2011 is initially designed for fine-grained classification. The resolution of all images in the three datasets is resized to  $84 \times 84$ . CIFAR-FS is a subset of CIFAR-100 and each image is resized to  $32 \times 32$ . See Appendix A.3 for more details on the four few-shot benchmark datasets.

**Implementation Details.** We adopt the episodic training procedure under 5-way 1-shot and 5-shot settings [38]. In each episode, 15 unlabeled query images per class are used for the training and testing phases. We apply Conv4 (filters:64) and ResNet12 (filters: [64, 160, 320, 640]) as the encoder. Our model is trained from scratch and uses the Adam optimizer with an initial learning rate  $10^{-3}$ . The hyperparameter  $\beta_j = 0.1$  for all experiments. Each mini-batch contains four episodes and we use a validation set to select the best-trained model. For all methods, we train 60,000 episodes for 1-shot and 40,000 episodes for 5-shot [4]. We use PyTorch with one NVIDIA Tesla V100 GPU to implement all experiments and report the *average accuracy (%)* with 95% *confidence interval* over the 10,000 randomly sampled testing episodes. For our proposed HTS framework, we consider ProtoNet as our FSL model unless otherwise stated, but note that our framework is broadly applicable to other meta-learning based FSL models. More implementation details are given in Appendix A.3.

**Pretext Tasks.** Following [16], as the entire input images during training is important for image classification, we choose two same pretext tasks: *rotation* and *color permutation*. We also give some subsets of the two tasks to meet the needs of experiments. The rotation task has **rotation1** ( $90^\circ$ ), **rotation2** ( $90^\circ$ ,  $180^\circ$ ), **rotation3** ( $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ) and **rotation4** ( $0^\circ$ ,  $90^\circ$ ,  $180^\circ$ ,  $270^\circ$ ). The color permutation task has **color\_perm1** (GBR), **color\_perm2** (GBR, BRG), **color\_perm3** (RGB, GBR, BRG) and **color\_perm6** (RGB, GBR, BRG, GBR, GRB, BRG, BGR). Note that our method can use arbitrary pretext tasks as it adaptively learns efficient features and improves the performance of FSL task.

## 5.2 RQ1. Performance of Pretext Tasks in FSL

The experimental results in Fig. 3 illustrate the key motivation of this paper and answer RQ1. Intuitively, when we use the pretext tasks to generate multiple sets of augmented images, it can learn more knowledge and improve the performance of the FSL model, like [26]. However, we find that not arbitrary pretext task has good performance on downstream datasets, and we need to find suitable tasks

| Training → Testing                             | Method         | 1-shot     | 5-shot     |
|--|----------------|------------|------------|
| <i>miniImageNet</i> →<br>CUB-200-2011          | ProtoNet + HTS | 32.18±0.25 | 55.95±0.21 |
|  |                | 39.57±0.17 | 60.56±0.18 |
| <i>tieredImageNet</i> →<br>CUB-200-2011        | ProtoNet + HTS | 39.47±0.22 | 56.58±0.25 |
|  |                | 42.24±0.20 | 60.71±0.18 |
| <i>tieredImageNet</i> →<br><i>miniImageNet</i> | ProtoNet + HTS | 47.01±0.26 | 66.82±0.25 |
|  |                | 55.29±0.20 | 72.67±0.16 |

**Table 1.** Classification accuracy (%) results comparison with 95% confidence intervals for cross-domain evaluation with **rotation3**. More results see Appendix A.4.

| Method            | Backbone | Tricks           | <i>miniImageNet</i> |                   | CUB-200-2011      |                   | CIFAR-FS          |                   | <i>tieredImageNet</i> |                   |
|-------------------|----------|------------------|---------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-----------------------|-------------------|
|                   |          |                  | 1-shot              | 5-shot            | 1-shot            | 5-shot            | 1-shot            | 5-shot            | 1-shot                | 5-shot            |
| MAML [10]         | Conv4-32 | Fine-tuning      | 48.70±1.84          | 55.31±0.73        | 55.92±0.95        | 72.09±0.76        | 58.90±1.90        | 71.50±1.00        | 51.67±1.81            | 70.30±1.75        |
| PN [11]           | Conv4-64 | Train-SSL        | 53.63±0.43          | 71.70±0.36        | -                 | -                 | 64.69±0.32        | <b>80.82±0.24</b> | -                     | -                 |
| CC [11]           | Conv4-64 | Train-SSL        | 54.83±0.43          | 71.86±0.33        | -                 | -                 | 63.45±0.31        | 79.79±0.24        | -                     | -                 |
| closer [4]        | Conv4-64 | Train-DA         | 48.24±0.75          | 66.43±0.63        | 60.53±0.83        | 79.34±0.61        | -                 | -                 | -                     | -                 |
| CSS [1]           | Conv4-64 | Train-SSL        | 50.85±0.84          | 68.08±0.73        | 66.01±0.90        | <b>81.84±0.59</b> | 56.49±0.93        | 74.59±0.72        | -                     | -                 |
| SLA [16]          | Conv4-64 | Train-SSL        | 44.95±0.79          | 63.32±0.68        | 48.43±0.82        | 71.30±0.72        | 45.94±0.87        | 68.62±0.75        | -                     | -                 |
| PSST [5]          | Conv4-64 | Train-SSL        | 57.04±0.51          | 73.17±0.48        | -                 | -                 | 64.37±0.33        | 80.42±0.32        | -                     | -                 |
| CAN [13]          | Conv4-64 | Train-DA         | 52.04±0.00          | 65.54±0.00        | -                 | -                 | -                 | -                 | 52.45±0.00            | 66.81±0.00        |
| <b>HTS (ours)</b> | Conv4-64 | Train&test-SSL   | <b>58.96±0.18</b>   | <b>75.17±0.14</b> | <b>67.32±0.24</b> | 78.09±0.15        | <b>64.71±0.21</b> | 76.45±0.17        | <b>53.20±0.22</b>     | <b>72.38±0.19</b> |
| shot-Free [28]    | ResNet12 | Train-DA         | 59.04±0.43          | 77.64±0.39        | -                 | -                 | 69.20±0.40        | 84.70±0.40        | 66.87±0.43            | 82.64±0.39        |
| MetaOpt [18]      | ResNet12 | Train-DA         | 62.64±0.61          | 78.63±0.46        | -                 | -                 | 72.00±0.70        | 84.20±0.50        | 65.81±0.74            | 82.64±0.39        |
| Distill [37]      | ResNet12 | Train&test-DA&KD | 64.82±0.60          | 82.14±0.43        | -                 | -                 | -                 | -                 | <b>71.52±0.69</b>     | 86.03±0.49        |
| <b>HTS (ours)</b> | ResNet12 | Train&test-SSL   | <b>64.95±0.18</b>   | <b>83.89±0.15</b> | <b>72.88±0.22</b> | <b>85.32±0.13</b> | <b>73.95±0.22</b> | <b>85.36±0.14</b> | 68.38±0.23            | <b>86.34±0.18</b> |

**Table 2.** Accuracy (%) with **rotation3**. Best results are displayed in boldface. Train/test-SSL/-DA mean using data augmentation and self-supervised learning during training/testing. KD means knowledge distillation and & means “and” operation.

for different datasets. The results of ProtoNet+DA are significantly lower than the baseline (ProtoNet), which proves our concern that it is unreasonable for empirical risk minimization (ERM) to treat all samples equally (see Eq. (3)). For ProtoNet+SSL, the results are slightly higher than the baseline in most cases, which indicates that only using a shared encoder is not enough (see Eq. (4)).

### 5.3 RQ2. Performance of Pretext Tasks in HTS


To answer RQ2, we conduct experiments on both single domain and cross domain with comparison to few-shot learning methods using the benchmark datasets.

**Performance on single domain.** Tab. 2 reports the average classification accuracy. For fairness, we use the same backbone network to compare to state-of-the-art meta-learning based FSL methods. From Tab. 2, we have the following findings: (1) HTS improves the performance of the ProtoNet in most settings to achieve a new state of the arts. This is because that our framework has the advantage of modeling relationships among these images and adaptively learning augmented features via a gated selection aggregating component. This observation demonstrates our motivation and the effectiveness of our framework. (2) One observation worth highlighting is that HTS not only outperforms traditional meta-learning based methods, but also is superior to methods using pretext tasks under DA or SSL settings. (3) Compared to [37], the results further show that augmenting the query set can bring more benefits during the testing phase.

**Performance on cross domain.** In Tab. 1, we show the testing results on testing domains using the model trained on training domains. The setting is challenging due to the domain gap between training and testing datasets. The results clearly show that (1) HTS has significant performance in all cross-domain set-

| Methods<br>under 5-way | <i>miniImageNet</i>   |                       | CUB-200-2011          |                       | CIFAR-FS              |                      | <i>tieredImageNet</i> |                       |
|------------------------|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|----------------------|-----------------------|-----------------------|
|                        | 1-shot                | 5-shot                | 1-shot                | 5-shot                | 1-shot                | 5-shot               | 1-shot                | 5-shot                |
| ProtoNet               | 44.42                 | 64.24                 | 51.31                 | 70.77                 | 51.90                 | 69.95                | 49.35                 | 67.28                 |
| [33] +HTS_DA           | 57.39                 | 74.25                 | 66.88                 | 77.37                 | <b>65.01</b> (+13.11) | 75.23                | 52.92                 | 70.18                 |
| +HTS_SSL               | <b>58.96</b> (+14.54) | <b>75.17</b> (+10.93) | <b>67.32</b> (+16.01) | <b>78.09</b> (+7.32)  | 64.71                 | <b>76.45</b> (+6.50) | <b>53.20</b> (+3.85)  | <b>72.38</b> (+5.10)  |
| RelationNet            | 49.31                 | 66.30                 | 62.45                 | 76.11                 | 55.00                 | 69.30                | 54.48                 | 71.32                 |
| [35] +HTS_DA           | 52.78                 | 73.22                 | 65.63                 | <b>80.67</b> (+4.56)  | 57.68                 | 73.09                | 55.44                 | <b>72.88</b> (+1.56)  |
| +HTS_SSL               | <b>53.24</b> (+3.93)  | <b>73.57</b> (+7.27)  | <b>67.38</b> (+4.93)  | 79.00                 | <b>58.60</b> (+3.60)  | <b>73.15</b> (+3.85) | <b>56.09</b> (+0.61)  | 71.78                 |
| MatchingNet            | 48.14                 | 63.48                 | 61.16                 | 72.86                 | 53.00                 | 60.23                | 54.02                 | 70.11                 |
| [38] +HTS_DA           | <b>53.64</b> (+5.50)  | 63.67                 | 63.29                 | 73.17                 | <b>55.57</b> (+2.57)  | 62.58                | 56.39                 | 72.01                 |
| +HTS_SSL               | 52.29                 | <b>64.36</b> (+0.88)  | <b>63.54</b> (+4.93)  | <b>74.76</b> (+1.90)  | 55.18                 | <b>63.87</b> (+3.64) | <b>57.16</b> (+3.14)  | <b>72.60</b> (+2.49)  |
| GNN                    | 49.02                 | 63.50                 | 51.83                 | 63.69                 | 45.59                 | 65.62                | 43.56                 | 55.31                 |
| [32] +HTS_DA           | 59.06                 | 72.80                 | 61.69                 | 73.46                 | 52.35                 | 71.66                | 55.32                 | 69.48                 |
| +HTS_SSL               | <b>60.52</b> (+11.50) | <b>74.63</b> (+11.13) | <b>62.85</b> (+11.02) | <b>77.58</b> (+13.89) | <b>58.31</b> (+12.72) | <b>73.24</b> (+7.62) | <b>55.73</b> (+12.17) | <b>70.42</b> (+15.11) |

**Table 3.** Accuracy (%) by incorporating HTS (two-level trees) into each method with rotation3 and the Conv4-64. Best results are displayed in **boldface** and performance improvement in red text. Appendix A.2 shows the formulation of these methods.



**Fig. 4.** (a)-(c) represent the t-SNE of the five class features. (b) and (c) are our method.


tings and obtains consistent improvements. (2) When CUB-200-2011 is selected as the testing domain, the transferred performance of different training domains has a large difference. It is caused by the degrees of domain gap. From the last two rows, we find that all methods from *tieredImageNet* to *miniImageNet* have a large improvement compared with CUB-200-2011, because the two datasets come from the same database (ILSVRC-12), leading to a small domain gap.

**Performance with other meta-learning based methods.** To further verify the effectiveness of HTS, we embed it into four meta-learning based methods: ProtoNet [33], RelationNet [35], MatchingNet [38] and GNN [32]. Tab. 3 reports the accuracy and we have the following findings: (1) our HTS is flexible and can be combined with any meta-learning method, making the performance of these methods significantly improved in all datasets. (2) In terms of our method, the HTS\_SSL performs better than HTS\_DA in most cases, indicating that a single label space is not best for learning the knowledge carried by these pretext tasks.

**T-SNE visualization.** For the qualitative analysis, we also apply t-SNE [15] to visualize the embedding distribution obtained before and after being equipped with HTS in ProtoNet. As shown in Fig. 4, (b) and (c) represent the features obtained without and with using the TreeLSTM aggregator. The results show that (1) our HTS can learn more compact and separate clusters indicating that the learned representations are more discriminative. (2) Considering our method, (c) is better than (b), which once again proves the effectiveness of the aggregator.

#### 5.4 RQ3. Adaptive Selection Aggregation

One of the most important properties of our framework is that the learned augmented features are different for different pretext tasks or child nodes. Thus, in this subsection, we investigate if the proposed framework can adaptively learn different forgetting gates for different child nodes, which aims to answer RQ3.



**Fig. 5.** (a) is the forgetting gate for different child nodes (vertical axis) with the same raw image (horizontal axis). (b) is the correlations of different augmented images (vertical axis is rotation1, horizontal axis is color\_perm1), where the diagonal (or off-diagonal) line is based on the same raw (or different raw) images. (c)-(d) are classification accuracy (%) for different numbers of levels under ProtoNet with HTS.


We visualize the average forgetting value (*i.e.*,  $f_m$  in Eq. (6)) for each child node with the same raw image randomly sampled from CUB-200-2011. In Fig. 5 (a), we show the results of six raw images as we have similar observations on other images and datasets. The darkness of a step represents the probability that the step is selected for aggregation. We can find that (1) different types of child nodes exhibit different forgetting values based on the same raw image. (2) Different raw images focus on different child node features. These results further explain that the contributions of different pretext tasks to the raw images are not equal. In Fig. 5 (b), we also provide the correlation, *i.e.*, the cosine similarity of learned features of different types of augmentation based on the same raw image. We clearly observe the correlations between augmentation with the same raw image are large while the correlation between different raw images is small, which meets our expectation that multiple sets of augmentation generated based on the same raw image generally have similar relationships.

### 5.5 RQ4. Ablation Study

To answer RQ4, we conduct experiments to show the performance under different pretext tasks, child nodes and levels. More experiments see Appendix A.5.

**The effect of different pretext tasks.** To show the performance of different pretext tasks, we construct experiments based on the two-level trees (*i.e.* only using a pretext task) for each specific pretext task with the same number of child nodes on the *miniImageNet* and CUB-200-2011 under the 5-way 1-shot setting. Fig. 6 (a) and (b) show the results. Along the horizontal axis, the number 0 indicates the baseline (*i.e.*, ProtoNet) only using raw (FSL) images. The rotation1/2/3 and color\_perm1/2/3 are used for the numbers 1/2/3, respectively. We follow these findings: (1) all pretext tasks significantly improve the classification performance and outperform the baseline. These results indicate that our HTS does not need to manually define a suitable pretext task, and it can adaptively learn useful information. (2) The color permutation task brings a higher boost on *miniImageNet* and the rotation task has a good performance on CUB-200-2011. It is because for fine-grained image classification, rotation angles can carry affluent representations, but for a large of images, the color gain is higher.

**The effect of different numbers of child nodes.** We verify the performance of the same pretext task with different numbers of child nodes on the



**Fig. 6.** (a) and (b) indicate different pretext tasks with the same number of child nodes. (c) and (d) indicate different numbers of child nodes with same pretext task.

*miniImageNet* under the 5-way 1-shot setting. Fig. 6 (c) and (d) show the results. Along the horizontal axis, these integers are introduced in Sec. 5.1. From Fig. 6 (c)-(d), we have the following findings: (1) our proposed HTS method with the two-level trees outperforms the ProtoNet in terms of all settings, which demonstrates the effectiveness of our method to learn pretext tasks. (2) The performance with different numbers of child nodes is different. The reason is that the increase of child nodes is accompanied by large model complexity, leading to overfitting and affecting the generalization ability of the FSL model.

**The effect of different numbers of tree levels.** We further verify the performance of HTS with different numbers of levels on the four datasets under the 5-way 1-shot and 5-shot settings. Fig. 5 (c) and (d) show the results. Along the horizontal axis, the integer 1 indicates only using FSL images, 2 using [rotation3], 3 using [rotation3; color\_perm2], 4 using [rotation3; color\_perm2; color\_perm3] and 5 using [rotation3; color\_perm2; color\_perm3; rotation2]. We find that our HTS method with different tree levels outperforms the original baseline (ProtoNet) using only raw (FSL) images. The different number of tree levels bring model complexity and over-fitting problems, resulting in performance differences among different levels. To save computational time and memory, two-level trees are used, which achieves the new state-of-the-art performance.

## 6 Conclusion

In this paper, we study the problem of learning richer and transferable representations for few-shot image classification via pretext tasks. We propose a plug-in *hierarchical tree structure-aware (HTS)* method that constructs tree structures to address this problem. Specifically, we propose the hierarchical tree constructing component, which uses the edge information to model the relationships of different augmented and raw images, and uses the level to explore the knowledge among different pretext tasks. On the constructed tree structures, we also introduce the gated selection aggregating component, which uses the forgetting gate to adaptively select and aggregate the augmented features and improve the performance of FSL. Extensive experimental results demonstrate that our HTS can achieve a new state-of-the-art performance on four benchmark datasets.

**Acknowledgments:** This work was supported by the National Science and Technology Innovation 2030 - Major Project (Grant No. 2022ZD0208800), and NSFC General Program (Grant No. 62176215). We thank Dr. Zhitao Wang for helpful feedback and discussions.

## References

1. An, Y., Xue, H., Zhao, X., Zhang, L.: Conditional self-supervised learning for few-shot classification. In: International Joint Conference on Artificial Intelligence, IJCAI (2021)
2. Bateni, P., Goyal, R., Masrani, V., Wood, F., Sigal, L.: Improved few-shot visual classification. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. pp. 14481–14490 (2020)
3. Bertinetto, L., Henriques, J.F., Torr, P.H.S., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: International Conference on Learning Representations, ICLR (2019)
4. Chen, W., Liu, Y., Kira, Z., Wang, Y.F., Huang, J.: A closer look at few-shot classification. In: International Conference on Learning Representations, ICLR (2019)
5. Chen, Z., Ge, J., Zhan, H., Huang, S., Wang, D.: Pareto self-supervised training for few-shot learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2021)
6. Cubuk, E.D., Zoph, B., Mané, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation strategies from data. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2019)
7. Cui, W., Guo, Y.: Parameterless transductive feature re-representation for few-shot learning. In: International Conference on Machine Learning, ICML (2021)
8. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Li, F.: Imagenet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2009)
9. Feng, Z., Xu, C., Tao, D.: Self-supervised representation learning by rotation feature decoupling. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2019)
10. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning, ICML (2017)
11. Gidaris, S., Bursuc, A., Komodakis, N., Pérez, P., Cord, M.: Boosting few-shot visual learning with self-supervision. In: International Conference on Computer Vision, ICCV (2019)
12. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations, ICLR (2018)
13. Hou, R., Chang, H., Ma, B., Shan, S., Chen, X.: Cross attention network for few-shot classification. In: Advances in Neural Information Processing Systems, NeurIPS (2019)
14. Kang, D., Kwon, H., Min, J., Cho, M.: Relational embedding for few-shot classification. In: International Conference on Computer Vision, ICCV (2021)
15. Laurens, Maaten, V.D., Hinton, Geoffrey: Visualizing data using t-sne. *Journal of Machine Learning Research* (2008)
16. Lee, H., Hwang, S.J., Shin, J.: Self-supervised label augmentation via input transformations. In: International Conference on Machine Learning, ICML (2020)
17. Li, F., Fergus, R., Perona, P.: A bayesian approach to unsupervised one-shot learning of object categories. In: International Conference on Computer Vision, ICCV (2003)
18. Li, H., Eigen, D., Dodge, S., Zeiler, M., Wang, X.: Finding task-relevant features for few-shot learning by category traversal. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2019)



19. Li, W., Wang, L., Xu, J., Huo, J., Gao, Y., Luo, J.: Revisiting local descriptor based image-to-class measure for few-shot learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2019)
20. Li, W., Xu, J., Huo, J., Wang, L., Gao, Y., Luo, J.: Distribution consistency based covariance metric networks for few-shot learning. In: Association for the Advancement of Artificial Intelligence, AAAI (2019)
21. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few shot learning. CoRR **abs/1707.09835** (2017)
22. Liu, L., Hamilton, W.L., Long, G., Jiang, J., Larochelle, H.: A universal representation transformer layer for few-shot image classification. In: International Conference on Learning Representations, ICLR. OpenReview.net (2021)
23. Liu, L., Zhou, T., Long, G., Jiang, J., Zhang, C.: Learning to propagate for graph meta-learning. In: Advances in Neural Information Processing Systems, NeurIPS. pp. 1037–1048 (2019)
24. Liu, S., Davison, A.J., Johns, E.: Self-supervised generalisation with meta auxiliary learning. In: Advances in Neural Information Processing Systems, NeurIPS (2020)
25. Misra, I., van der Maaten, L.: Self-supervised learning of pretext-invariant representations. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2020)
26. Ni, R., Goldblum, M., Sharaf, A., Kong, K., Goldstein, T.: Data augmentation for meta-learning. In: International Conference on Machine Learning, ICML (2021)
27. Qiao, S., Liu, C., Shen, W., Yuille, A.L.: Few-shot image recognition by predicting parameters from activations. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2018)
28. Ravichandran, A., Bhotika, R., Soatto, S.: Few-shot learning with embedded class models and shot-free meta training. In: International Conference on Computer Vision , ICCV (2019)
29. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: International Conference on Learning Representations, ICLR (2018)
30. Requeima, J., Gordon, J., Bronskill, J., Nowozin, S., Turner, R.E.: Fast and flexible multi-task classification using conditional neural adaptive processes. In: Advances in Neural Information Processing Systems, NeurIPS. pp. 7957–7968 (2019)
31. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: International Conference on Learning Representations, ICLR (2019)
32. Satorras, V.G., Estrach, J.B.: Few-shot learning with graph neural networks. In: International Conference on Learning Representations, ICLR (2018)
33. Snell, J., Swersky, K., Zemel, R.S.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems, NeurIPS (2017)
34. Su, J., Maji, S., Hariharan, B.: When does self-supervision improve few-shot learning? In: European Conference on Computer Vision, ECCV (2020)
35. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H.S., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2018)
36. Tai, K.S., Socher, R., Manning, C.D.: Improved semantic representations from tree-structured long short-term memory networks. In: Association for Computational Linguistics, ACL (2015)
37. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: A good embedding is all you need? In: European Conference on Computer Vision, ECCV (2020)

38. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems*, NeurIPS (2016)
39. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. California Institute of Technology (CNS-TR-2011-001) (2011)
40. Yao, H., Wei, Y., Huang, J., Li, Z.: Hierarchically structured meta-learning. In: *International Conference on Machine Learning, ICML. Proceedings of Machine Learning Research*, vol. 97, pp. 7045–7054. PMLR (2019)
41. Ye, H., Hu, H., Zhan, D., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR* (2020)
42. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 12200–12210 (2020)
43. Zhang, M., Zhang, J., Lu, Z., Xiang, T., Ding, M., Huang, S.: Iept: Instance-level and episode-level pretext tasks for few-shot learning. In: *International Conference on Learning Representations, ICLR* (2020)
44. Zhang, M., Huang, S., Wang, D.: Domain generalized few-shot image classification via meta regularization network. In: *ICASSP*. pp. 3748–3752 (2022)
45. Zhang, M., Wang, D., Gai, S.: Knowledge distillation for model-agnostic meta-learning. In: *European Conference on Artificial Intelligence, ECAI* (2020)