

Meta-Transfer Learning through Hard Tasks

Qianru Sun*, Yaoyao Liu*, Zhaozheng Chen, Tat-Seng Chua, and Bernt Schiele, *Fellow, IEEE*

Abstract—Meta-learning has been proposed as a framework to address the challenging few-shot learning setting. The key idea is to leverage a large number of similar few-shot tasks in order to learn how to adapt a base-learner to a new task for which only a few labeled samples are available. As deep neural networks (DNNs) tend to overfit using a few samples only, typical meta-learning models use shallow neural networks, thus limiting its effectiveness. In order to achieve top performance, some recent works tried to use the DNNs pre-trained on large-scale datasets but mostly in straight-forward manners, e.g., (1) taking their weights as a warm start of meta-training, and (2) freezing their convolutional layers as the feature extractor of base-learners. In this paper, we propose a novel approach called **meta-transfer learning (MTL)** which **learns to transfer the weights of a deep NN** for few-shot learning tasks. Specifically, *meta* refers to training multiple tasks, and *transfer* is achieved by learning scaling and shifting functions of DNN weights for each task. In addition, we introduce the **hard task (HT) meta-batch** scheme as an effective learning curriculum that further boosts the learning efficiency of MTL. We conduct few-shot learning experiments and report top performance for five-class few-shot recognition tasks on three challenging benchmarks: minImageNet, tieredImageNet and Fewshot-CIFAR100 (FC100). Extensive comparisons to related works validate that our MTL approach trained with the proposed HT meta-batch scheme achieves top performance. An ablation study also shows that both components contribute to fast convergence and high accuracy.

Compared to the conference version of the paper [1], this version additionally presents (1) the analysis of using different DNN architectures, e.g., ResNet-12, ResNet-18 and ResNet-25; (2) the new MTL variants that adapt our scaling and shifting functions to the classical supervised and the state-of-the-art semi-supervised meta-learners, and achieve performance improvements consistently; (3) the discussion of new related works since the conference version; and (4) the results on the larger and more challenging benchmark – tieredImageNet [2].



1 INTRODUCTION

ALTHOUGH deep learning systems have achieved great performance when sufficient amounts of labeled data are available [3]–[5], there has been growing interest in reducing the required amount of data. Few-shot learning tasks have been defined for this purpose. The aim is to learn new concepts from a handful of training examples, e.g. from 1 or 5 training images [1], [6], [7]. Humans tend to be highly effective in this context, often grasping the essential connection between new concepts and their own knowledge and experience, but it remains challenging for machine learning models. For instance on the CIFAR-100 dataset, a classification model trained in the fully supervised mode achieves 76% accuracy for the 100-class setting [8], while the best-performing 1-shot model achieves only 45% in average for the much simpler 5-class setting [1]. On the other hand, in many real-world applications we are lacking large-scale training data, as e.g. in the medical domain. It is thus desirable to improve machine learning models in order to handle few-shot settings.

Basically, the nature of few-shot learning with very scarce training data makes it difficult to train powerful machine learning models for new concepts. People explore a variety of methods in order to overcome this. A straight forward idea is to increase the amount of available data by data augmentation techniques [10]. Several methods proposed to learn a data generator e.g. con-

ditioned on Gaussian noises [11]–[13] or object attributes [14]. However, this data generator often under-performs when trained on few-shot data, which has been investigated by [15]. An alternative is to merge data from multiple tasks which, however, is often ineffective due to high variances of the data across tasks [13].

In contrast to data augmentation methods, meta-learning is a task-level learning method [16]–[18]. It aims to transfer experience from similar few-shot learning tasks [1], [7], [19]–[24]. Related methods follow a unified training process that contains two loops. The inner-loop learns a base-learner for an individual task, and the outer-loop then uses the validation performance of the learned base-learner to optimize the meta-learner. A state-of-the-art representative method named Model-Agnostic Meta-Learning (MAML) learns to search for the optimal initialization state to fast adapt a base-learner to a new task [7]. Its task-agnostic property makes it possible to generalize to few-shot supervised/semi-supervised learning as well as unsupervised reinforcement learning [7], [19], [20], [23]–[25]. However, in our view, there are two main limitations of this type of approaches limiting their effectiveness: i) these methods usually require a large number of similar tasks for meta-training which is costly; and ii) each task is typically modeled by a low-complexity base-learner, such as a shallow neural network (SNN), to avoid model overfitting to few-shot training data, thus being unable to deploy deeper and more powerful network architectures. For example, for the miniImageNet dataset [26], MAML uses a *shallow* CNN with only 4 CONV layers and its optimal performance was obtained by learning on 240k tasks (60k iterations in total and each meta-batch contains 4 tasks).

In this paper, we propose a novel meta-learning method called **meta-transfer learning (MTL)** leveraging the advantages of both transfer learning and meta-learning (see conceptual comparison of related methods in the upper block of Figure 1). In a nutshell,

- *Qianru Sun is the corresponding author. qianrusun@smu.edu.sg. She and Zhaozheng Chen are with the School of Information Systems, Singapore Management University. Yaoyao Liu is with the School of Electrical and Information Engineering, Tianjin University. This work was done during his internship supervised by Qianru Sun. Tat-Seng Chua is with the School of Computing, National University of Singapore. Bernt Schiele is with the Department of Computer Vision and Machine Learning, Max-Planck Institute for Informatics.*
- * indicates equal contribution.

Manuscript received on Sep 29th, 2019.

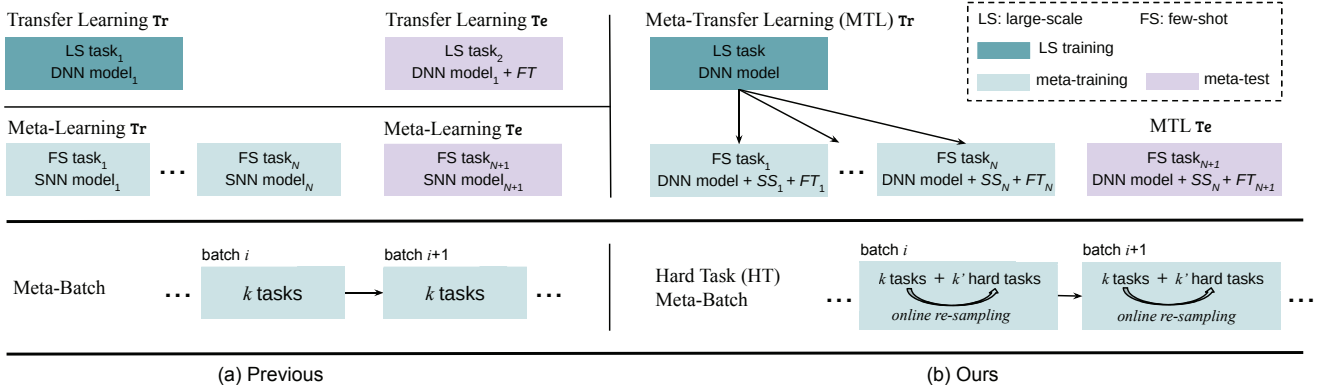


Fig. 1. Meta-transfer learning (MTL) is our meta-learning paradigm and hard task (HT) meta-batch is our training strategy. The upper blocks show the differences between MTL and related methods – transfer-learning [9] and meta-learning [7]. The bottom blocks compare the proposed HT meta-batch with the conventional meta-batch [7]. Note that here the *FT* stands for fine-tuning a classifier. *SS* represents the *Scaling* and *Shifting* operations in our MTL method.

MTL is a novel learning method that helps deep neural networks (DNNs) converge faster while reducing the probability to overfit when training on few labeled data only. In particular, “transfer” means that DNN weights trained on large-scale data can be used in other tasks by two light-weight neuron operations: *Scaling* and *Shifting* (*SS*), i.e. $\alpha X + \beta$. “Meta” means that the parameters of these operations can be viewed as hyper-parameters trained on few-shot learning tasks [21], [27], [28]. First, large-scale trained DNN weights offer a good initialization, enabling fast convergence of MTL with fewer tasks, e.g., only $8k$ tasks for miniImageNet [26], 30 times fewer than MAML [7]. Second, light-weight operations on DNN neurons have less parameters to learn, e.g., less than $\frac{2}{49}$ if considering neurons of size 7×7 ($\frac{1}{49}$ for α and $< \frac{1}{49}$ for β), reducing the chance of overfitting to few-shot data. Third, these operations keep those trained DNN weights unchanged, and thus avoid the problem of “catastrophic forgetting” which means forgetting general patterns when adapting to a specific task [29], [30]. Finally, these operations are conducted on the convolutional layers mostly working for image feature extraction, thus can generalize well to a variety of few-shot learning models, e.g., MAML [7], MatchingNet [26], ProtoNet [31], RelationNet [32], and Baseline++ [33].

The second main contribution of this paper is an effective meta-training curriculum. Curriculum learning [34] and hard negative mining [35] both suggest that faster convergence and stronger performance can be achieved by better arrangement of training data. Inspired by these ideas, we design our **hard task (HT) meta-batch** strategy to offer a challenging but effective learning curriculum. As shown in the bottom rows of Figure 1, a conventional meta-batch contains a number of random tasks [7], but our HT meta-batch online re-samples harder ones according to past failure tasks with lowest validation accuracy.

Our overall contribution is thus three-fold: i) we propose a novel **MTL** method that learns to transfer large-scale pre-trained DNN weights for solving few-shot learning tasks; ii) we propose a novel **HT meta-batch** learning strategy that forces meta-transfer to “grow faster and stronger through hardship”; and iii) we conduct extensive experiments on three few-shot learning benchmarks, namely miniImageNet [26], tieredImageNet [2] and Fewshot-CIFAR100 (FC100) [36], and achieve the state-of-the-art performance. **Compared to the conference version** of the paper [1], this version additionally presents (1) the results of

using different DNN architectures, e.g., ResNet-12, ResNet-18 and ResNet-25, (2) new MTL variants that combine our scaling and shifting functions with the classical supervised and the state-of-the-art semi-supervised meta-learners models to achieve top performance in respective scenarios, compared to post-conference works, (3) discussion of new related works since the conference version, and (4) results on the larger and more challenging benchmark – tieredImageNet [2]. Our Tensorflow and Pytorch codes are open-sourced at github.com/yaoyao-liu/meta-transfer-learning.

2 RELATED WORK

Research literature on few-shot learning exhibits great diversity, spanning from data augmentation [11]–[14] to supervised meta-learning [18], [37]. In this paper, we focus on the meta-learning based methods most relevant to ours and compared to in the experiments. Besides, we borrow the idea of transfer learning when leveraging the large-scale pre-training step in prior to meta-transfer. For task sampling, our HT meta-batch scheme is related to curriculum learning and hard negative sampling methods.

Meta-learning. We can divide meta-learning methods into three categories. 1) *Metric learning* methods learn a similarity space in which learning is particularly efficient for few-shot training examples. Examples of distance metrics include cosine similarity [26], [33], Euclidean distance to the prototypical representation of a class [31], CNN-based relation module [32], ridge regression based [38], and graph model based [39], [40]. Some recent works also tried to generate task-specific feature representation for few-shot episodes based on metric learning, like [41], [42] 2) *Memory network* methods learn to store “experience” when learning seen tasks and then generalize it to unseen tasks. The key idea is to design a model specifically for fast learning with a few training steps. A family of model architectures use external memory storage include Neural Turing Machines [43], Meta Networks [27], Neural Attentive Learner (SNAIL) [44], and Task Dependent Adaptive Metric (TADAM) [36]. For test, general meta memory and specific task information are combined to make predictions in neural networks. 3) *Gradient descent* based meta-learning methods intend for adjusting the optimization algorithm so that the model can converge within a small number of optimization steps (with a few examples). The optimization algorithm can be explicitly modeled with two learning loops that outer-loop has a *meta-learner* that

learns to adapt an inner-loop *base-learner* (to few-shot examples) through different tasks. For example, Ravi *et al.* [45] introduced a method that compresses the base-learners’ parameter space in an LSTM meta-learner. Rusu *et al.* [25] designed a classifier generator as the meta-learner which output parameters for each specific base-learning task. Finn *et al.* [7] proposed a meta-learner called MAML that learns to effectively initialize a base-learner for a new task. Lee *et al.* [46] presented a meta-learning approach with convex base-learners for few-shot tasks. Other related works in this category include Hierarchical Bayesian model [20], Bilevel Programming [21], and GAN based meta model [23].

Among them, MAML is a fairly general optimization algorithm, compatible with any model that learns through gradient descent. Its meta-learner optimization is done by gradient descent using the validation loss of the base-learner. It is closely related to our MTL. An important difference is that MTL leverages transfer learning and benefits from referencing neuron knowledge in pre-trained deep nets. Although MAML can start from a pre-trained network, its element-wise fine-tuning makes it hard to learn deep nets without overfitting (validated in our experiments).

Transfer learning Transfer learning or knowledge transfer has the goal to transfer the information of trained models to solve unknown tasks, thereby reducing the effort to collect new training data. *What* and *how* to transfer are key issues to be addressed. Different methods are applied to different source-target domains and bridge different transfer knowledge [9], [47]–[50]. For deep models, a powerful transfer method is adapting a pre-trained model for a new task, often called *fine-tuning* (*FT*). Models pre-trained on large-scale datasets have proven to generalize better than randomly initialized ones [51]. Another popular transfer method is taking pre-trained networks as backbone and adding high-level functions, e.g. for object detection [52] and image segmentation [53], [54]. Besides, the knowledge to transfer can be from multi-modal category models, e.g. the word embedding models used for zero-shot learning [14], [55] and trained attribute models used for social relationship recognition [50].

In this paper, our meta-transfer learning leverages the idea of transferring pre-trained weights and our model meta-learns how to effectively transfer. The large-scale trained DNN weights are *what* to transfer, and the operations of *Scaling* and *Shifting* indicate *how* to transfer. Some few-shot learning methods have been proposed to utilize pre-trained DNNs [25], [44], [56]–[58]. Typically, DNN weights are either fixed for feature extraction or simply fine-tuned for each task, while we learn a meta-transfer learner through all tasks, which is different in terms of the underlying learning paradigm. More importantly, our approach can generalize to existing few-shot learning models whose image features are extracted from DNNs on different architectures, for which we conduct extensive experiments in Section 5.

Curriculum learning & Hard sample mining Curriculum learning was proposed by Bengio *et al.* [34] and is popular for multi-task learning [59]–[61]. They showed that instead of observing samples at random it is better to organize samples in a meaningful way so that fast convergence, effective learning and better generalization can be achieved. Kumar *et al.* [62] introduced an iterative self-paced learning algorithm where each iteration simultaneously selects easy samples and learns a new parameter vector. Intuitively, the curriculum is determined by the pupils abilities rather than being fixed by a teacher. Pentina *et al.* [63] use adaptive SVM classifiers to evaluate task difficulty for later organization. Most

recently, Jiang *et al.* [64] designed a MentorNet that provides a “curriculum”, i.e., sample weighting scheme, for StudentNet to focus on the labels which are probably correct. The trained MentorNet can be directly applied for the training of StudentNet on a new dataset. Differently, our MTL method does task difficulty evaluation online at the phase of test in each task, without needing any auxiliary model.

Hard sample mining was proposed by Shrivastava *et al.* [35] for object detection with DNNs. It treats image proposals overlapped with ground truth (i.e. causing more confusions) as hard negative samples. Training on more confusing data enables the detection model to achieve higher robustness and better performance [65]–[67]. Inspired by this, we sample harder tasks online and make our MTL learner “grow faster and stronger through more hardness”. In our experiments, we show that this can be generalized to different architectures with different meta-training operations, i.e. *SS* and *FT*, referring to Figure 6.

3 PRELIMINARY

In this section, we briefly introduce the unified episodic formulation in meta-learning, following related works [7], [25], [26], [36], [45]. Then, we introduce the task-level data denotations used at two phases, i.e., meta-train and meta-test.

Meta-learning has an episodic formulation which was proposed for tackling few-shot tasks first in [26]. It is different from traditional image classification, in three aspects: (1) the main phases are not train and test but meta-train and meta-test, each of which includes training and testing; (2) the samples in meta-train and meta-test are not datapoints but episodes, and each episode is a few-shot classification task; and (3) the objective is not classifying unseen datapoints but to fast adapt the meta-learned experience or knowledge to the learning of a new few-shot classification task.

The denotations of two phases, meta-train and meta-test, are as follows. A meta-train example is a classification task \mathcal{T} sampled from a distribution $p(\mathcal{T})$. \mathcal{T} is called episode, including a training split $\mathcal{T}^{(tr)}$ to optimize the base-learner, i.e., the classifiers in our model, and a test split $\mathcal{T}^{(te)}$ to optimize the meta-learner, i.e., the scaling and shifting parameters in our model. In particular, meta-train aims to learn from a number of episodes $\{\mathcal{T}\}$ sampled from $p(\mathcal{T})$. An unseen task \mathcal{T}_{unseen} in meta-test will start from that experience of the meta-learner and adapt the base-learner. The final evaluation is done by testing a set of unseen datapoints in $\mathcal{T}_{unseen}^{(te)}$.

Meta-train phase. This phase aims to learn a meta-learner from multiple episodes. In each episode, meta-training has a two-stage optimization. Stage-1 is called base-learning, where the cross-entropy loss is used to optimize the parameters of the base-learner. Stage-2 contains a feed-forward test on episode test datapoints. The test loss is used to optimize the parameters of the meta-learner. Specifically, given an episode $\mathcal{T} \in p(\mathcal{T})$, the base-learner $\theta_{\mathcal{T}}$ is learned from episode training data $\mathcal{T}^{(tr)}$ and its corresponding loss $\mathcal{L}_{\mathcal{T}}(\theta_{\mathcal{T}}, \mathcal{T}^{(tr)})$. After optimizing this loss, the base-learner has parameters $\tilde{\theta}_{\mathcal{T}}$. Then, the meta-learner is updated using test loss $\mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathcal{T}}, \mathcal{T}^{(te)})$. After meta-training on all episodes, the meta-learner is optimized by test losses $\{\mathcal{L}_{\mathcal{T}}(\tilde{\theta}_{\mathcal{T}}, \mathcal{T}^{(te)})\}_{\mathcal{T} \in p(\mathcal{T})}$. Therefore, the number of meta-learner updates equals to the number of episodes.

Meta-test phase. This phase aims to test the performance of the trained meta-learner for fast adaptation to unseen task. Given \mathcal{T}_{unseen} , the meta-learner $\tilde{\theta}_{\mathcal{T}}$ teaches the base-learner $\theta_{\mathcal{T}_{unseen}}$ to

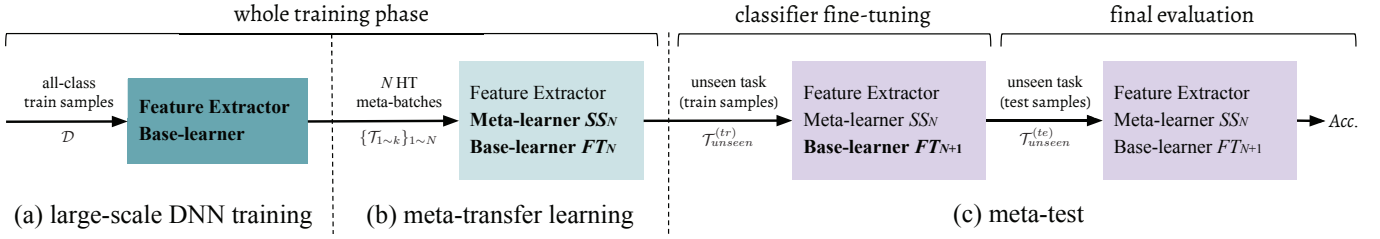


Fig. 2. The pipeline of our proposed few-shot learning method, including three phases: (a) DNN training on large-scale data, i.e. using all training datapoints (Section 4.1); (b) Meta-transfer learning (MTL) that learns the parameters of *Scaling* and *Shifting* (*SS*), based on the pre-trained feature extractor (Section 4.2). Learning is scheduled by the proposed HT meta-batch (Section 4.4); and (c) meta-test is done for an unseen task which consists of a base-learner (classifier) *Fine-Tuning* (*FT*) stage and a final evaluation stage, described in the last paragraph in Section 3. Input data are along with arrows. Modules with names in bold get updated at corresponding phases.

adapt to the objective of \mathcal{T}_{unseen} by some means, e.g. through initialization [7]. Then, the test result on $\mathcal{T}_{unseen}^{(te)}$ is used to evaluate the meta-learning approach. If there are multiple unseen tasks $\{\mathcal{T}_{unseen}^{(te)}\}$, the average result on $\{\mathcal{T}_{unseen}^{(te)}\}$ will be the final evaluation.

4 METHODOLOGY

As shown in Figure 2, our method consists of three phases. First, we train a DNN on large-scale data, e.g. on miniImageNet with 64 classes and 600 samples per class [26], and then fix the low-level layers as Feature Extractor (Section 4.1). Second, in the meta-transfer learning phase, MTL learns the *Scaling* and *Shifting* (*SS*) parameters for the Feature Extractor neurons, enabling fast adaptation to few-shot tasks (Section 4.2). For improving the overall learning, we use our HT meta-batch strategy (Section 4.4). The training steps are detailed in Algorithm 1 and Algorithm 2 in Section 4.5. Thirdly, the typical meta-test phase is performed, as introduced in Section 3, and the details are given in Algorithm 3.

4.1 DNN training on large-scale data

This phase is similar to the classic pre-training stage as, e.g., pre-training on Imagenet for object recognition [68]. Here, we do not consider data/domain adaptation from other datasets, and pre-train on readily available data of few-shot learning benchmarks, allowing for fair comparison with other few-shot learning methods. Specifically, for a particular few-shot dataset, we merge all-class data \mathcal{D} for pre-training. For instance, for miniImageNet [26], there are totally 64 classes in the training split \mathcal{D} and each class contains 600 samples, which we use to pre-train a 64-class classifier.

We first randomly initialize a feature extractor Θ (e.g. CONV layers in ResNets [4]) and a classifier θ (e.g. the last FC layer in ResNets [4]), and then optimize them by gradient descent as follows,

$$[\Theta; \theta] =: [\Theta; \theta] - \alpha \nabla \mathcal{L}_{\mathcal{D}}([\Theta; \theta]), \quad (1)$$

where \mathcal{L} denotes the following empirical loss,

$$\mathcal{L}_{\mathcal{D}}([\Theta; \theta]) = \frac{1}{|\mathcal{D}|} \sum_{(x,y) \in \mathcal{D}} l(f_{[\Theta; \theta]}(x), y), \quad (2)$$

e.g. cross-entropy loss, and α denotes the learning rate. In this phase, the feature extractor Θ is learned. It will be frozen in the following meta-training and meta-test phases, as shown in Figure 2. The learned classifier θ will be discarded, because subsequent few-shot tasks contain different classification objectives, e.g. 5-class instead of 64-class classification for miniImageNet [26].

4.2 Meta-transfer learning (MTL)

As shown in Figure 2(b), our proposed meta-transfer learning (MTL) method optimizes the meta operations *Scaling* and *Shifting* (*SS*) through HT meta-batch training (Section 4.4). Figure 4 visualizes the difference of updating through *SS* and *Fine-Tuning* (*FT*). *SS* operations, denoted as Φ_{S_1} and Φ_{S_2} , do not change the frozen neuron weights of Θ during learning, while *FT* updates the complete Θ .

In the following, we expand the details of *SS* operations, corresponding to Figure 2 (b) and Figure 3. Given a task \mathcal{T} , the loss of $\mathcal{T}^{(tr)}$ is used to optimize the current base-learner (classifier) θ' by gradient descent:

$$\theta' \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_{\mathcal{T}^{(tr)}}([\Theta; \theta], \Phi_{S_{\{1,2\}}}), \quad (3)$$

which is different to Eq. 1, as we do not update Θ . Note that here θ is different to the one from the previous phase, the large-scale classifier θ in Eq. 1. This θ concerns only a few classes, e.g. 5 classes, to classify each time in a novel few-shot setting. θ' corresponds to a temporal classifier only working in the current task, initialized by the θ optimized for the previous task (see Eq. 5).

Φ_{S_1} is initialized by ones and Φ_{S_2} by zeros. Then, they are optimized by the test loss of $\mathcal{T}^{(te)}$ as follows,

$$\Phi_{S_i} =: \Phi_{S_i} - \gamma \nabla_{\Phi_{S_i}} \mathcal{L}_{\mathcal{T}^{(te)}}([\Theta; \theta'], \Phi_{S_{\{1,2\}}}), i = 1, 2. \quad (4)$$

In this step, θ is updated with the same learning rate γ as in Eq. 4,

$$\theta =: \theta - \gamma \nabla_{\theta} \mathcal{L}_{\mathcal{T}^{(te)}}([\Theta; \theta'], \Phi_{S_{\{1,2\}}}). \quad (5)$$

Re-linking to Eq. 3, we note that the above θ' comes from the last epoch of base-learning on $\mathcal{T}^{(tr)}$.

Next, we describe how we apply $\Phi_{S_{\{1,2\}}}$ to the frozen neurons as shown in Figure 4(b). Given the trained Θ , for its l -th layer containing K neurons, we have K pairs of parameters, respectively as weight and bias, denoted as $\{(W_{i,k}, b_{i,k})\}$. Note that the neuron location l, k will be omitted for readability. Based on MTL, we learn K pairs of scalars $\{\Phi_{S_{\{1,2\}}}\}$. Assuming X is input, we apply $\{\Phi_{S_{\{1,2\}}}\}$ to (W, b) as

$$SS(X; W, b; \Phi_{S_{\{1,2\}}}) = (W \odot \Phi_{S_1})X + (b + \Phi_{S_2}), \quad (6)$$

where \odot denotes the element-wise multiplication.

Taking Figure 4(b) as an example of a single 3×3 filter, after *SS* operations, this filter is scaled by Φ_{S_1} then the feature maps after convolutions are shifted by Φ_{S_2} in addition to the original bias b . Detailed steps of *SS* are given in Algorithm 2 in Section 4.5.

Figure 4(a) shows a typical parameter-level *FT* operation, which is in the meta optimization phase of our related work

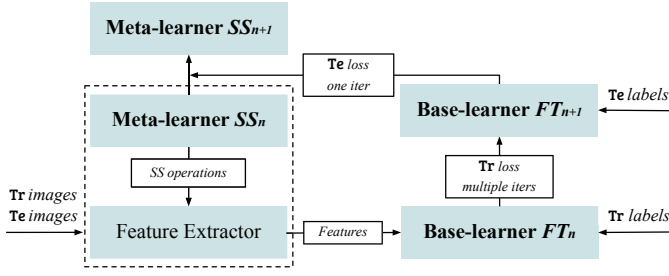


Fig. 3. The computation flow on one meta-train task. It includes the update on the parameters of base-learner and meta-learner. Base-learner is initialized by the weights meta-learned from previous tasks, and its update is the conventional fine-tuning (FT). Note that this FT works in base-learning, thus is different with the meta-operation FT used in the outer-loops of MAML [7]. When this update finishes, the average loss of the test data are computed to update the MTL meta-learner, i.e. SS parameters.

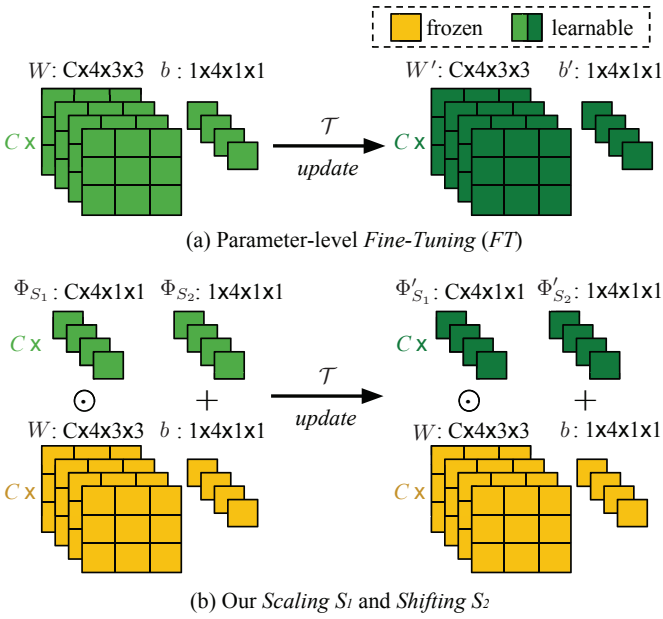


Fig. 4. Two kinds of meta operations on pre-trained weights. (a) Parameter-level *Fine-Tuning* (FT) is a conventional meta-train operation used in related works such as MAML [7], ProtoNets [31] and RelationNets [32]. Its update works for all neuron parameters, W and b . (b) Our neuron-level *Scaling* and *Shifting* (SS) operations in MTL. They reduce the number of learning parameters and avoid overfitting problems. In addition, they keep large-scale trained parameters (in yellow) frozen, preventing “catastrophic forgetting” [29], [30].

MAML [7]. It is obvious that FT updates the complete values of W and b , and has a large number of parameters, and our SS reduces this number to below $\frac{2}{9}$ in the example of the figure.

In summary, SS can benefit MTL in three aspects. 1) It starts from a strong initialization based on a large-scale trained DNN, yielding fast convergence for MTL. 2) It does not change DNN weights, thereby avoiding the problem of “catastrophic forgetting” [29], [30] when learning specific tasks in MTL. 3) It is light-weight, reducing the chance of overfitting of MTL in few-shot scenarios.

4.3 Generalized MTL

As introduced, key components of our meta-transfer learning include a large-scale pre-training and two meta operations – *Scaling* and *Shifting* pre-trained weights. As shown in Figure 3,

SS operations make the effect on feature extraction, i.e., backbone networks, for learning a new task. They are thus independent of the classifier architecture which is typically the “head” of the end-to-end network.

While, most metric-based and optimization-based meta learning models have same backbone “body” for feature extraction but different classification “heads” for classification. Metric-based methods use nearest neighbor classifiers, e.g., MatchingNets [26], ProtoNets [31], and RelationNets [32]. Optimization-based methods learn through backpropagation of unrolled gradients of multiple iterations of the classifier, and different methods have different designs for classifier, e.g., Cosine-similarity based classifier [33], [69], Softmax classifier [1], [7], [19], and SVM [46].

In the conference paper, we use Softmax classifier which is a simple linear FC layer in deep nets. In this journal version, we generalize meta-transfer learning method to other models. In total, we implement five classical models including Cosine-similarity based model [33], [69], MatchingNets [26], ProtoNets [31], RelationNets [32] and our previous Softmax model [1], [7], [19]. In experiments, we show that under the same conditions of network architectures and hyperparameters, using our deep pre-trained model with SS operations brings consistent improvements, in terms of classification accuracy as well as learning speed, over previously reported results.

4.4 Hard task (HT) meta-batch

In this section, we introduce a method to schedule hard tasks in meta-training batches. The conventional meta-batch is composed of randomly sampled tasks, where the randomness implies random difficulties [7]. In our meta-training pipeline, we intentionally pick up failure cases in each task and re-compose their data to be harder tasks for adverse re-training. The task flow is shown in Figure 5. We aim to force our meta-learner to “grow up through hardness”.

Pipeline. Given a (M -class, N -shot) task \mathcal{T} , a meta-batch $\{\mathcal{T}_{1\sim k}\}$ contains two splits, $\mathcal{T}^{(tr)}$ and $\mathcal{T}^{(te)}$, for base-learning and test, respectively. As shown in Algorithm 2 line 2-5, base-learner is optimized by the loss of $\mathcal{T}^{(tr)}$ (in multiple epochs). SS parameters are then optimized by the loss of $\mathcal{T}^{(te)}$ once. During the loss computation on $\mathcal{T}^{(te)}$, we can also get the recognition accuracy for M classes. Then, we choose the lowest accuracy Acc_{m^*} to determine the most difficult class m^* (also called failure class) in the current task.

After obtaining all failure classes $\{m^*\}$ from $\{\mathcal{T}_{1\sim k}\}$ in current meta-batch (k is the batch size), we re-sample tasks from the data indexed by $\{m^*\}$. Specifically, we assume $p(\mathcal{T}|\{m^*\})$ is the task distribution, we sample a “harder” task $\mathcal{T}^{hard} \in p(\mathcal{T}|\{m^*\})$. Two important details are given below.

Choosing hard class m^* . We choose the failure class m^* from each task by ranking the class-level accuracies instead of fixing a threshold. In a dynamic online setting as ours, it is more sensible to choose the hardest cases based on ranking rather than fixing a threshold ahead of time.

Two methods of hard tasking using $\{m^*\}$. Chosen $\{m^*\}$, we can re-sample tasks \mathcal{T}^{hard} by (1) directly using the samples of m^* -th class in the current task \mathcal{T} , or (2) indirectly using the index m^* to sample new samples of that class. In fact, setting (2) considers to include more data variance of m^* -th class and it works better than setting (1) in general.

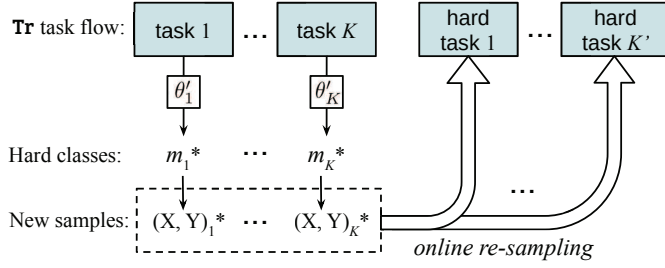


Fig. 5. The computation flow of online hard task sampling. During an HT meta-batch phase, the meta-training first goes through K random tasks then continues on re-sampled K' hard tasks.

4.5 Algorithm

The training steps of our MTL approach is provided in Algorithm 1 and Algorithm 2, and the meta-test steps in Algorithm 3.

Algorithm 1 summarizes the training process of two main stages: large-scale DNN training (line 1-5) and meta-transfer learning (line 6-19). HT meta-batch re-sampling and continuous training phases are shown in line 13-18, for which the failure classes are returned by Algorithm 2, see line 11.

Algorithm 1: MTL with HT meta-batch strategy

Input: Task distribution $p(\mathcal{T})$ and corresponding dataset \mathcal{D} , learning rates α, β and γ
Output: Feature extractor Θ , base learner θ , *Scaling* and *Shifting* parameters $\Phi_{S_{\{1,2\}}}$

- 1 Randomly initialize Θ and θ ;
- 2 **for** *samples* in \mathcal{D} **do**
- 3 Evaluate $\mathcal{L}_{\mathcal{D}}([\Theta; \theta])$ by Eq. 2;
- 4 Optimize Θ and θ by Eq. 1;
- 5 **end**
- 6 Initialize Φ_{S_1} by ones, initialize Φ_{S_2} by zeros; Reset θ for few-shot tasks; Randomly initialize θ ; Initialize $\{m^*\}$ as an empty set.
- 7 **for** *iterations in meta-training* **do**
- 8 Randomly sample a batch of tasks $\{\mathcal{T}_{1 \sim K}\} \in p(\mathcal{T})$;
- 9 **for** k from 1 to K **do**
- 10 Optimize $\Phi_{S_{\{1,2\}}}$ and θ with \mathcal{T}_k by Algorithm 2;
- 11 Get the returned the m^* -th class, then add it to a hard class set $\{m^*\}$;
- 12 **end**
- 13 Sample hard tasks $\{\mathcal{T}^{hard}\} \subseteq p(\mathcal{T}|\{m^*\})$;
- 14 **for** k from 1 to K' **do**
- 15 Sample task $\mathcal{T}_k^{hard} \in \{\mathcal{T}^{hard}\}$;
- 16 Optimize $\Phi_{S_{\{1,2\}}}$ and θ with \mathcal{T}_k^{hard} by Algorithm 2;
- 17 **end**
- 18 Empty $\{m^*\}$.
- 19 **end**

Algorithm 2 presents an entire learning epoch on a single task. Base-learning steps in the episode training split are given in line 2-5. The meta-level update by the test loss is shown in line 6. In line 7-11, the recognition rates of all test classes are computed and returned to Algorithm 1 (line 11) for hard task sampling.

Algorithm 3 presents the evaluation process on unseen tasks. The average accuracy Acc obtained in this Algorithm is used as the final evaluation of our method.

Algorithm 2: Detail learning steps of a training task

Input: Task \mathcal{T} , learning rates β and γ , feature extractor Θ , base learner θ , *Scaling* and *Shifting* parameters $\Phi_{S_{\{1,2\}}}$
Output: Base learner θ , *Scaling* and *Shifting* parameters $\Phi_{S_{\{1,2\}}}$, the worst classified class- m in \mathcal{T}

- 1 Sample training datapoints $\mathcal{T}^{(tr)}$ and test datapoints $\mathcal{T}^{(te)}$ from \mathcal{T} ;
- 2 **for** *samples* in $\mathcal{T}^{(tr)}$ **do**
- 3 Evaluate $\mathcal{L}_{\mathcal{T}^{(tr)}}$;
- 4 Optimize θ' by Eq. 3;
- 5 **end**
- 6 Optimize $\Phi_{S_{\{1,2\}}}$ and θ by Eq. 4 and Eq. 5;
- 7 **for** $m \in \{1 \sim M\}$ **do**
- 8 Classify samples of m -th class in $\mathcal{T}^{(te)}$;
- 9 Compute Acc_m ;
- 10 **end**
- 11 Return the m^* -th class with the lowest accuracy Acc_m^* .

Algorithm 3: Meta test on the unseen task

Input: Unseen task \mathcal{T}_{un} , base learning rate β and meta learning rate γ , feature extractor Θ , base learner θ , *Scaling* and *Shifting* parameters $\Phi_{S_{\{1,2\}}}$
Output: Average accuracy Acc

- 1 Sample training datapoints $\mathcal{T}_{un}^{(tr)}$ and test datapoints $\mathcal{T}_{un}^{(te)}$ from \mathcal{T}_{un} ;
- 2 **for** *samples* in $\mathcal{T}_{un}^{(tr)}$ **do**
- 3 Evaluate $\mathcal{L}_{\mathcal{T}_{un}^{(tr)}}$;
- 4 Optimize θ' by Eq. 3;
- 5 **end**
- 6 **for** *samples* in $\mathcal{T}_{un}^{(te)}$ **do**
- 7 Predict class labels by $[\Theta; \theta']$;
- 8 Obtain the classification result;
- 9 **end**
- 10 Compute the accuracy Acc_m of m -th class;
- 11 Compute the mean accuracy Acc of M classes.

5 EXPERIMENTS

We evaluate the proposed MTL and HT meta-batch in terms of few-shot recognition accuracy and model convergence speed. Comparing to our conference version [1], this paper contains additional experiments on (i) a larger and more challenging dataset – tieredImageNet [2] which have been widely adopted in recent related works [25], [40], [42], [46]; (ii) two deeper architectures – Resnet-18 and ResNet-25 as the Feature Extractor Θ ; and (iii) supervised and semi-supervised MTL variances in which we implement SS operations on both classical and state-of-the-art meta-learners, i.e. ProtoNets [31], MatchingNets [26], RelationNets [32], and the meta-learning version of Baseline++ (superior to original Baseline++ [33]), as well as the state-of-the-art semi-supervised few-shot classification models, i.e. Soft Masked k -Means [2] and TPN [40].

Below we describe the datasets we evaluate on and detailed settings, followed by the comparisons to state-of-the-art methods, extensive validations on multiple MTL variances, and an ablation study regarding SS operations and HT meta-batch.

5.1 Datasets

We conduct few-shot learning experiments on three benchmarks, miniImageNet [26], tieredImageNet [2] and Fewshot-CIFAR100 (FC100) [36]. miniImageNet is the most widely used in related works [7], [20], [21], [45], [70], and the later ones are more recently published with a larger scale and a more challenging setting, i.e., lower image resolution and stricter training-test splits.

miniImageNet [26]. It was proposed especially for few-shot learning evaluation. Its complexity is high due to the use of ImageNet images, but it requires less resource and infrastructure than running on the full ImageNet dataset [68]. In total, there are 100 classes with 600 samples of 84×84 color images per class. These 100 classes are divided into 64, 16, and 20 classes respectively for sampling tasks for meta-training, meta-validation and meta-test, following related works [7], [20], [21], [45], [70].

tieredImageNet [2]. Compared to miniImageNet, it is a larger subset of ImageNet with 608 classes (779, 165 images) grouped into 34 super-class nodes. These nodes are partitioned into 20, 6, and 8 disjoint sets respectively for meta-training, validation, and test. The corresponding sub-classes are used to build the classification tasks in each of which the 5 sub-classes are randomly sampled. As argued in [2], this super-class based training-test split results in a more challenging and realistic regime with meta-test and validation tasks that are less similar to meta-training tasks.

Fewshot-CIFAR100 (FC100) [36]. This dataset is based on the popular object classification dataset CIFAR100 [71]. Its training-test splits are also based on super-classes [36]. In total, it contains 100 object classes (600 images per class) belonging to 20 super-classes. Meta-training data are from 60 classes belonging to 12 super-classes. Meta-validation and meta-test sets contain 20 classes belonging to 4 super-classes, respectively. Comparing to the ImageNet subsets above, FC100 offers a more challenging scenario with lower image resolution, i.e. each sample is a 32×32 color image.

5.2 Implementation details

Task sampling. We use the same task sampling method as related works [7], on all datasets. Specifically, (1) we consider the 5-class classification, (2) during meta-training, we sample 5-class, 1-shot (or 5-shot) episodes to contain 1 (or 5) samples for train episode and 15 (uniform) samples for episode test, and (3) during meta-validation and meta-test, we sample 5-class, 1-shot (or 5-shot) episodes to contain 1 (or 5) samples for train episode and 1 (uniform) sample for episode test. Note that in some related works, e.g. [36], 32 samples are used for episode test on 5-shot tasks. Using such a larger number of test samples results in the lower standard variance of recognition accuracies.

In total, we sample at most $20k$ tasks ($10k$ meta-batches) for meta-training (same for the cases *w/* and *w/o* HT meta-batch), and sample 600 random tasks for both meta-validation and meta-test [7]. Note that we choose the trained models which have the highest meta-validation accuracies, for meta-test.

Network architectures. We present the details of network architectures for Feature Extractor parameters Θ , MTL meta-learner with *Scaling* and *Shifting* parameters Φ_{S_1} , Φ_{S_2} , and MTL base-learner (classifier) parameters θ . For Θ , in our conference version [1], we used ResNet-12 and 4CONV which have been commonly used in previous works [7], [26], [36], [44]–[46], [70]. In this journal version, we implement two deeper architectures – ResNet-18 and ResNet-25, which have been adopted in newly

published related works [42], [57], [72], and we achieve the top performance using ResNet-25.

Specifically, **4CONV** consists of 4 layers with 3×3 convolutions and 32 filters, followed by batch normalization (BN) [73], a ReLU nonlinearity, and 2×2 max-pooling. MTL only works with the following deep nets. **ResNet-12** contains 4 residual blocks and each block has 3 CONV layers with 3×3 kernels. At the end of each residual block, a 2×2 max-pooling layer is applied. The number of filters starts from 64 and is doubled every next block. Following 4 blocks, there is a mean-pooling layer to compress the output feature maps to a feature embedding. **ResNet-18** contains 4 residual blocks and each block has 4 CONV layers with 3×3 kernels. The number of filters starts from 64 and is doubled every next block. Before the residual blocks, there is one additional CONV layer with 64 filter and 3×3 kernels at the beginning of the network. The residual blocks are followed by an average pooling layer. The ResNet-18 backbone we use exactly follow [4] except that the last FC layer is removed. **ResNet-25** is exactly the same as the released code of [57], [74]. Three residual blocks are used after an initial convolutional layer. Each block has 4 CONV layers with 3×3 kernels. The number of filters starts from 160 and is doubled every next block. After a global average pooling layer, it leads to a 640-dim embedding. For fair comparison with important baseline methods such as MAML [7] and ProtoNets [31] which are based on 4CONV nets, we implement the experiments for them using the same ResNet-12 and ResNet-18 architectures.

For the architecture of Φ_{S_1} and Φ_{S_2} , actually, they are generated according to the architecture of Θ , as introduced in Section 4.2. For example, when using ResNet-25 in MTL, Φ_{S_1} and Φ_{S_2} also have 25 layers, respectively.

For the architecture of θ (the parameters of the base-learner), we empirically find that in our cases a single FC layer (as θ) is faster to train and more effective for classification than multiple layers, taking the most popular dataset miniImageNet as an example. Results are given in Table 1, in which we can see the performance drop when changing this θ to multiple layers.

Pre-training stage. For the phase of DNN training on large-scale data (Section 4.1), the model is trained by Adam optimizer [75]. The learning rate is initialized as 0.001, and decays to its half every $5k$ iterations until it is lower than 0.0001. We set the keep probability of the dropout as 0.9 and batch-size as 64. The pre-training stops after $10k$ iterations. Note that for hyperparameter selection, we randomly choose 550 samples each class as the training set, and the rest as validation. After the grid search for hyperparameters, we fix them and mix up all samples (64 classes, 600 samples each class) to do the final pre-training. Image samples in these steps are augmented by horizontal flipping.

Meta-training stage. This is a task-level training in which the base-learning in one task considers a training step for optimizing base-learner, followed by a validation step for optimizing meta-learner. The base-learner θ is optimized by batch gradient descent with the learning rate of 0.01. It is updated with 20 and 60 epochs respectively for 1-shot and 5-shot tasks on the miniImageNet and tieredImageNet datasets, and 20 epochs for all tasks on the FC100 dataset. Specially when using ResNet-25, we use 100 epochs for all tasks on all datasets. The meta-learner, i.e. the parameters of the *SS* operations, is optimized by Adam optimizer [75]. Its learning rate is initialized as 0.001, and decays to the half every $1k$ iterations until 0.0001. The size of meta-batch is set to 2 (tasks) due to the memory limit.

HT meta-batch. Hard tasks are sampled every time after running 10 meta-batches, i.e., the failure classes used for sampling hard tasks are from 20 tasks as each meta-batch contains 2 tasks. The number of hard tasks is selected for different settings by validation: 10 and 4 hard tasks respectively for the 1-shot and 5-shot experiments, on the miniImageNet and tieredImageNet datasets; and respectively 20 and 10 hard tasks for the 1-shot, 5-shot experiments, on the FC100 dataset.

Base-learning	Dim. of θ	miniImageNet	
		1-shot	5-shot
θ (2 FC layers)	512, 5	59.1 \pm 1.9	70.7 \pm 0.9
θ (3 FC layers)	1024, 512, 5	56.2 \pm 1.8	68.7 \pm 0.9
Θ, θ	5	59.6 \pm 1.8	71.6 \pm 0.9
θ (Ours)	5	60.6 \pm 1.9	74.3 \pm 0.8

TABLE 1

The 5-way, 1-shot and 5-shot classification accuracy (%) on miniImageNet, when using different architectures for the base-learner (i.e., the classifier θ).

5.3 Ablation setting

In order to show the effectiveness of our approach, we carefully design several ablative settings: two baselines without meta-learning but more classic learning, named as *update**, four baselines of *Fine-Tuning (FT)* on different numbers of parameters in the outer-loop based on MAML [7], named as *FT**, and two *SS* variants on smaller numbers of parameters, named as *SS**. Table 5 shows the results in these settings, for which we simply use the classical architecture (ResNet-12) containing 4 residual blocks named $\Theta 1 \sim \Theta 4$ and an FC layer θ (classifier). The bullet names used in the Table are explained as follows.

update $[\Theta; \theta]$ (or θ). There is no meta-training phase. During test phase, each task has its whole model $[\Theta; \theta]$ (or the classifier θ) updated on $\mathcal{T}^{(tr)}$, and then tested on $\mathcal{T}^{(te)}$.

FT θ ($[\Theta 4; \theta]$ or $[\Theta 3; \Theta 4; \theta]$ or θ). These are straight-forward ways to reduce the quantity of meta-learned parameters. For example, “ $[\Theta 3; \Theta 4; \theta]$ ” does not update the the first two residual blocks which encode the low-level image features. Specially, “ θ ” means only the classifier parameters are updated during meta-training.

SS $[\Theta 4; \theta]$ (or $[\Theta 3; \Theta 4; \theta]$ or θ). During the meta-training, a limited number of *SS* parameters are defined and used. Low-level residual blocks simply use the pre-trained weights without meta-level update.

5.4 Comparison with related works

Table 2 presents the overall comparisons with related works, on the miniImageNet, tieredImageNet, and FC100 datasets. Note that these numbers are the meta-test results of the meta-trained models which have the highest meta-validation accuracies. Specifically, on the miniImageNet, models on 1-shot and 5-shot are meta-trained for $6k$ and $10k$ iterations, respectively. On the tieredImageNet, iterations for 1-shot and 5-shot are at $8k$ and $10k$, respectively. On the FC100, iterations are all at $3k$.

miniImageNet. In the first block of Table 2, we can see that the ResNet-25 based MTL model learned in HT meta-batch scheme achieves the best few-shot classification performance, i.e. 80.9%,

for (5-class, 5-shot). It tackles the 1-shot tasks with an accuracy of 64.3%, comparable to the state-of-the-art result (69.13%) reported by LGM-Net [41] whose 5-shot result is 9.7% lower than ours. Regarding the network architecture, we can see that models using deeper ones, i.e. ResNet-12, ResNet-18, and ResNet-25, outperforms the 4CONV based models by rather large margins, e.g. 4CONV models have the best 1-shot result with 55.51% [40] which is 8.8% lower than our 64.3%. This demonstrates our contribution of utilizing deeper neural networks to better tackle the few-shot classification problems.

tieredImageNet. In the second block of Table 2, we give the results on the large-scale dataset – tieredImageNet. Since this dataset is newly proposed [2], its results using classical models [7], [31], [32] were reported in recent papers [2], [40]. From the table, we can see that the ResNet-25 MTL outperforms others, e.g. it achieves around a 4% margin over the top-performed CTM [42] on 1-shot tasks. An interesting observation is that on this larger and more challenging dataset, our deeper version of MTL (with ResNet-25) outperforms the shallower one (with ResNet-12) by 6.4% on 1-shot, which is the double of the improvement (3.2%) obtained on miniImageNet. This demonstrates that our idea of using deeper neural networks shows to be more promising for handling more difficult few-shot learning scenarios.

FC100. In the last block of Table 2, we show the results on the FC100. We report the results of TADAM [36] and MetaOptNet [46] using the numbers given in original papers, and obtain the results of classical methods – MAML [7], MAML++ [24], RelationNets [32] and MatchineNets [26] by implementing their open-sourced code with our deep networks. From these results, we can see that MTL consistently outperforms the original MAML and the improved version – MAML++ by large margins, e.g. around 7% in the 1-shot cases. It also surpasses both TADAM and MetaOptNet by about 5%.

From the overview of Table 2, we can conclude that our large-scale network pre-training and transferring on deeper CNNs significantly boost the few-shot learning performance. In this journal version, we add experiments on the tieredImageNet dataset whose meta-train and meta-test are clearly split according to super-classes. It is interesting to observe that our performance improvements on this more challenging dataset is clearly more than those on miniImageNet, and this is more obvious on the 1-shot tasks for which our best model achieves 3.6% higher than the new state-of-the-art model – CTM [42].

5.5 Generalization ability of MTL

The scaling and shifting (*SS*) operations in our proposed MTL approach work on pre-trained convolutional neurons thus are easy to be applied to other CNNs base few-shot learning models.

Table 3 shows the results of implementing *SS* operations on pre-trained deep models and using classical few-shot learning methods, namely ProtoNets [31], MatchingNets [26], RelationNets [32], MAML [7] (with a single FC layer as the base-learner [1]), and MAML [7] (with a cosine distance classifier as the base-learner [26], [33]), under fair data split settings. In the original versions of those methods, *FT* is the meta-level operation and 4CONV is the uniform architecture. For comparison, we, therefore, implement the results of using *FT* on deeper networks – ResNet-12 and ResNet-18. In Table 3, we have three column results. “miniImageNet(tieredPre)” denotes that the model is pre-trained on tieredImageNet and its weights are then meta-transferred to the learning of few-shot models on miniImageNet

Few-shot Learning Method		Backbone	miniImageNet (test)	
			1-shot	5-shot
<i>Data augmentation</i>	Adv. ResNet, [11]	WRN-40 (pre)	55.2	69.6
	Delta-encoder, [12]	VGG-16 (pre)	58.7	73.6
<i>Metric learning</i>	MatchingNets, [26]	4 CONV	43.44 ± 0.77	55.31 ± 0.73
	ProtoNets, [31]	4 CONV	49.42 ± 0.78	68.20 ± 0.66
	RelationNets, [32]	4 CONV	50.44 ± 0.82	65.32 ± 0.70
	Graph neural network, [39]	4 CONV	50.33 ± 0.36	66.41 ± 0.63
	Ridge regression, [38]	4 CONV	51.9 ± 0.2	68.7 ± 0.2
	TransductiveProp, [40]	4 CONV	55.51	69.86
<i>Memory network</i>	Meta Networks, [27]	5 CONV	49.21 ± 0.96	–
	SNAIL, [44]	ResNet-12 (pre) [◊]	55.71 ± 0.99	68.88 ± 0.92
	TADAM, [36]	ResNet-12 (pre) [†]	58.5 ± 0.3	76.7 ± 0.3
<i>Gradient descent</i>	MAML, [7]	4 CONV	48.70 ± 1.75	63.11 ± 0.92
	Meta-LSTM, [45]	4 CONV	43.56 ± 0.84	60.60 ± 0.71
	Hierarchical Bayes, [20]	4 CONV	49.40 ± 1.83	–
	Bilevel Programming, [21]	ResNet-12 [◊]	50.54 ± 0.85	64.53 ± 0.68
	MetaGAN, [23]	ResNet-12	52.71 ± 0.64	68.63 ± 0.67
	adaResNet, [70]	ResNet-12 [‡]	56.88 ± 0.62	71.94 ± 0.57
	MetaOptNet, [46]	ResNet-12	62.64 ± 0.35	78.63 ± 0.68
	LEO, [25]	WRN-28-10 (pre)	61.67 ± 0.08	77.59 ± 0.12
	LGM-Net, [41]	MetaNet+4CONV	69.13 ± 0.35	71.18 ± 0.68
CTM, [42]	ResNet-18 (pre)	64.12 ± 0.82	80.51 ± 0.13	
Ours	<i>FT</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	59.1 ± 1.9	73.1 ± 0.9
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	61.2 ± 1.8	75.5 ± 0.8
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-18 (pre)	61.7 ± 1.8	75.6 ± 0.9
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-25 (pre)	64.3 ± 1.7	80.9 ± 0.8
Few-shot Learning Method		Backbone	tieredImageNet (test)	
			1-shot	5-shot
<i>Metric learning</i>	ProtoNets, [31] (by [2])	4 CONV	53.31 ± 0.89	72.69 ± 0.74
	RelationNets, [32] (by [40])	4 CONV	54.48 ± 0.93	71.32 ± 0.78
	TransductiveProp, [40]	4 CONV	57.41 ± 0.94	71.55 ± 0.74
<i>Gradient descent</i>	MAML, [7] (by [40])	ResNet-12	51.67 ± 1.81	70.30 ± 0.08
	LEO, [25]	WRN-28-10 (pre)	66.33 ± 0.05	81.44 ± 0.09
	CTM, [42]	ResNet-18 (pre)	68.41 ± 0.39	84.28 ± 1.73
Ours	<i>FT</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	64.8 ± 1.9	78.4 ± 0.9
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	65.6 ± 1.8	80.8 ± 0.8
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-18 (pre)	67.5 ± 1.9	82.5 ± 0.7
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-25 (pre)	72.0 ± 1.8	85.1 ± 0.8
Few-shot Learning Method		Backbone	FC100 (test)	
			1-shot	5-shot
<i>Gradient descent</i>	MAML, [7] (by us)	4 CONV	38.1 ± 1.7	50.4 ± 1.0
	MAML++, [24] (by us)	4 CONV	38.7 ± 0.4	52.9 ± 0.4
	MetaOptNet [46]	4 CONV	41.1 ± 0.6	55.5 ± 0.6
<i>Memory network</i>	TADAM, [36]	ResNet-12 (pre) [†]	40.1 ± 0.4	56.1 ± 0.4
<i>Metric learning</i>	MatchingNets, [26] (by us)	ResNet-25 (pre)	44.2 ± 1.8	58.0 ± 0.8
	RelationNets, [32] (by us)	ResNet-25 (pre)	41.1 ± 1.8	58.6 ± 0.8
Ours	<i>FT</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	41.8 ± 1.9	55.1 ± 0.9
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-12 (pre)	45.1 ± 1.9	57.6 ± 0.9
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-18 (pre)	45.5 ± 1.9	57.6 ± 1.0
	<i>SS</i> [Θ; θ], HT meta-batch	ResNet-25 (pre)	46.0 ± 1.7	61.3 ± 0.8

[◊]Additional 2 convolutional layers [‡]Additional 1 convolutional layer [†]Additional 72 fully connected layers

TABLE 2

The 5-way, 1-shot and 5-shot classification accuracy (%) on miniImageNet, tieredImageNet, and FC100 datasets. “pre” means including our pre-training step with all training datapoints. “by [*]” means the results were reported in [*]. “by us” means the results are from our implementation of open-sourced code. Note that the standard variance is affected by the number of episode test samples, and our sample splits are the same with MAML [7].

Backbone	Method	Operation	miniImageNet		miniImageNet (tieredPre)		tieredImageNet	
			1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
ResNet-12	ProtoNets [31]	SS	56.7 ± 1.9	72.0 ± 0.9	62.0 ± 1.9	77.9 ± 1.0	62.2 ± 2.1	78.1 ± 0.9
		FT	55.2 ± 1.9	70.8 ± 0.9	57.2 ± 1.9	75.9 ± 0.9	54.9 ± 2.0	73.0 ± 1.0
	MatchingNets [26]	SS	58.1 ± 1.8	66.9 ± 0.9	63.6 ± 1.7	73.2 ± 0.9	64.5 ± 1.9	73.9 ± 0.9
		FT	57.4 ± 1.7	67.5 ± 0.8	61.1 ± 1.8	72.6 ± 0.8	62.4 ± 1.8	73.5 ± 0.8
	RelationNets [32]	SS	57.2 ± 1.8	71.1 ± 0.9	61.5 ± 1.8	74.9 ± 0.9	65.6 ± 1.9	77.5 ± 0.9
		FT	56.0 ± 1.8	69.0 ± 0.8	58.9 ± 1.8	72.0 ± 0.8	62.2 ± 1.8	76.0 ± 0.9
	MTL (FC) MAML [7] (FC)	SS	60.6 ± 1.9	74.3 ± 0.8	65.7 ± 1.8	78.4 ± 0.8	65.6 ± 1.7	78.7 ± 0.9
		FT	58.3 ± 1.9	71.6 ± 0.9	61.6 ± 1.9	73.5 ± 0.8	62.0 ± 1.8	70.6 ± 0.9
	MTL (Cosine) MAML [7] (Cosine)	SS	58.2 ± 1.8	74.6 ± 0.8	66.1 ± 1.8	79.7 ± 0.9	67.1 ± 1.8	80.0 ± 0.8
		FT	59.8 ± 1.8	72.8 ± 0.9	59.9 ± 1.9	76.5 ± 0.7	65.1 ± 1.9	78.2 ± 0.8
ResNet-18	ProtoNets [31]	SS	55.6 ± 1.9	70.7 ± 0.9	63.0 ± w	78.3 ± 1.0	58.6 ± 2.0	77.8 ± 1.0
		FT	55.5 ± 1.8	70.4 ± 0.8	59.2 ± 1.8	74.6 ± 0.8	58.4 ± 1.7	75.0 ± 0.8
	MatchingNets [26]	SS	57.2 ± 1.9	65.9 ± 0.9	62.3 ± 1.8	74.0 ± 0.9	64.3 ± 1.9	75.5 ± 0.9
		FT	56.9 ± 1.7	68.4 ± 0.8	62.2 ± 1.8	72.0 ± 0.8	62.9 ± 1.8	75.9 ± 0.8
	RelationNets [32]	SS	57.6 ± 0.8	71.1 ± 0.8	60.7 ± 1.9	74.6 ± 1.9	61.9 ± 1.8	77.8 ± 0.9
		FT	51.5 ± 1.8	64.4 ± 0.8	46.8 ± 1.8	66.4 ± 0.8	58.4 ± 1.8	71.8 ± 0.8
	MTL (FC) MAML [7] (FC)	SS	60.8 ± 1.9	74.5 ± 0.8	66.5 ± 1.8	80.2 ± 0.8	67.0 ± 1.8	80.6 ± 0.8
		FT	56.8 ± 1.9	65.4 ± 0.7	50.9 ± 1.8	68.8 ± 0.8	50.5 ± 1.8	71.8 ± 0.8
	MTL (Cosine) MAML [7] (Cosine)	SS	59.4 ± 1.9	74.3 ± 0.8	66.7 ± 1.8	81.8 ± 0.8	66.6 ± 1.8	82.1 ± 0.8
		FT	58.1 ± 1.9	69.9 ± 0.8	62.1 ± 1.9	77.7 ± 0.8	64.8 ± 1.9	79.5 ± 0.8

TABLE 3

The 5-way, 1-shot and 5-shot classification accuracy (%) on miniImageNet and tieredImageNet datasets. “(tieredPre)” means the pre-training stage is finished on the tieredImageNet. We implement the public code of related methods [7], [26], [31]–[33] in our framework by which we are able to conduct different meta Operations, i.e. *FT* and *SS*. Note that (1) cosine classifiers have been used in MatchingNets [31] and Baseline++ [33] for few-shot classification; and (2) MAML in this table is not exactly same with original MAML [7], as it works on deep neural networks and does not update convolutional layers during base-training.

	miniImageNet		tieredImageNet		miniImageNet w/D		tieredImageNet w/D	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
Masked Soft k-Means [2]	50.4 ± 0.3	64.4 ± 0.2	52.4 ± 0.4	69.9 ± 0.2	49.0 ± 0.3	63.0 ± 0.1	51.4 ± 0.4	69.1 ± 0.3
Masked Soft k-Means w/ MTL	58.2 ± 1.8	71.9 ± 0.8	65.3 ± 0.9	79.8 ± 0.8	56.8 ± 1.7	71.1 ± 0.8	63.6 ± 1.8	79.2 ± 0.8
TPN [40]	52.8 ± 0.3	66.4 ± 0.2	55.7 ± 0.3	71.0 ± 0.2	50.4 ± 0.8	64.9 ± 0.7	53.5 ± 0.9	69.9 ± 0.8
TPN w/ MTL	59.3 ± 1.7	71.9 ± 0.8	67.4 ± 1.8	80.7 ± 0.8	58.7 ± 1.7	70.6 ± 0.8	67.2 ± 1.7	80.5 ± 0.9

TABLE 4

Semi-supervised 5-way, 1-shot and 5-shot classification accuracy (%) on miniImageNet and tieredImageNet. “meta-batch” and “ResNet-12 (pre)” are used. “w/D” means additionally including the unlabeled data from 3 distracting classes (5 unlabeled samples per class) that are **excluded** in the “5-way” classes of the task [2], [40].

tasks. We can see that in all settings, (1) the best performance is always achieved by our proposed *SS* operations, e.g., MTL (FC, ResNet-12) outperforms MAML (FC, ResNet-12) by 3.6% and 8.1% on tieredImageNet 1-shot and 5-shot, respectively.; and (2) each classical method using *SS* gets consistent improvements its original *FT* version, e.g., RelationNets [32] (ResNet-18) gains 6.1% and 6.7% on miniImageNet 1-shot and 5-shot, respectively.

In addition, we verify the generalization ability of our MTL to the state-of-the-art semi-supervised few-shot learning (SSFSL) methods [2], [40]. To this end, we conduct experiments in the uniform SSFSL settings with 5 unlabeled samples per class added to the episode training set, following [2], [40]. Our results on the miniImageNet and tieredImageNet datasets are presented in Table 4. Note that “w/D” indicates the more challenging setting with 3 distracting classes added to the unlabeled data of the task. In each block of Table 4, we can clearly see that the methods

“w/ MTL” obtain consistent performance improvements over the original methods in both normal and challenging SSFSL settings by quit large margins, e.g. 13.7% on tieredImageNet w/D 1-shot. This validates the effectiveness of our method for tackling SSFSL problems.

5.6 Ablation study

Table 5 shows the extensive results of MTL-related ablative methods, on the miniImageNet and FC100 datasets. Figure 6 demonstrates the performance gap between w/ and w/o HT meta-batch in terms of accuracy and converging speed. Extensive comparisons are conducted on three datasets, three network architectures and two meta-level operations.

MTL vs. No meta-learning. Table 5 shows the results of *No meta-learning* methods on the top block. Compared to these, our approach achieves significantly better performance, e.g. the largest

	miniImageNet		miniImageNet (tieredPre)		FC100		FC100 (tieredPre)	
	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
<i>update</i> $[\Theta; \theta]$	45.3 \pm 1.9	64.6 \pm 0.9	54.4 \pm 1.8	73.7 \pm 0.8	38.4 \pm 1.8	52.6 \pm 0.9	37.6 \pm 1.9	52.7 \pm 0.9
<i>update</i> θ	50.0 \pm 1.8	66.7 \pm 0.9	51.3 \pm 1.8	70.3 \pm 0.8	39.3 \pm 1.9	51.8 \pm 0.9	38.3 \pm 1.8	52.9 \pm 1.0
<i>FT</i> θ	55.9 \pm 1.9	71.4 \pm 0.9	61.6 \pm 1.8	73.5 \pm 0.9	41.6 \pm 1.9	54.9 \pm 1.0	40.4 \pm 1.9	54.7 \pm 0.9
<i>FT</i> $[\Theta 4; \theta]$	57.2 \pm 1.8	71.6 \pm 0.8	62.3 \pm 1.8	73.9 \pm 0.9	40.9 \pm 1.8	54.3 \pm 1.0	41.2 \pm 1.8	53.6 \pm 1.0
<i>FT</i> $[\Theta 3, \Theta 4; \theta]$	58.1 \pm 1.8	70.9 \pm 0.8	63.0 \pm 1.8	74.8 \pm 0.9	41.5 \pm 1.8	53.7 \pm 0.9	40.7 \pm 1.9	53.8 \pm 0.9
<i>FT</i> $[\Theta; \theta]$	58.3 \pm 1.8	71.6 \pm 0.8	63.2 \pm 1.9	75.7 \pm 0.8	41.6 \pm 1.9	54.4 \pm 1.0	41.1 \pm 1.9	54.5 \pm 0.9
<i>SS</i> $[\Theta 4; \theta]$	59.2 \pm 1.8	73.1 \pm 0.9	64.0 \pm 1.8	76.9 \pm 0.8	42.4 \pm 1.9	55.1 \pm 1.0	42.7 \pm 1.9	55.9 \pm 1.0
<i>SS</i> $[\Theta 3, \Theta 4; \theta]$	59.4 \pm 1.8	73.4 \pm 0.8	64.5 \pm 1.8	77.2 \pm 0.8	42.5 \pm 1.9	54.5 \pm 1.0	43.4 \pm 1.8	56.4 \pm 1.0
<i>SS</i> $[\Theta; \theta]$ (Ours)	60.6 \pm 1.8	74.3 \pm 0.8	65.7 \pm 1.8	78.4 \pm 0.9	43.6 \pm 1.9	55.4 \pm 1.0	43.5 \pm 1.9	57.1 \pm 1.0

TABLE 5

The 5-way, 1-shot and 5-shot classification accuracy (%) using ablative models, on two datasets. “meta-batch” and “ResNet-12 (pre)” are used. “(tieredPre)” means the pre-training stage is finished on the tieredImageNet.

margins on miniImageNet are 10.6% for 1-shot and 7.6% for 5-shot. This validates the effectiveness of our meta-learning method for tackling few-shot learning problems. Between two *No meta-learning* methods, we can see that updating both feature extractor Θ and classifier θ is inferior to updating θ only (Θ is pre-trained), e.g. around 5% reduction on miniImageNet 1-shot. One reason is that in few-shot settings, there are too many parameters to optimize with little data. This supports our motivation to learn only θ during base-learning.

***SS* $[\Theta; \theta]$ works better than light-weight *FT* variants.** Table 5 shows that our approach with *SS* $[\Theta; \theta]$ achieves the best performances for all few-shot settings. In principle, *SS* meta-learns a smaller set of transferring parameters on $[\Theta; \theta]$ than *FT*. People may argue that *FT* is weaker because it learns a larger set of initialization parameters, whose quantity is equal to the size of $[\Theta; \theta]$, causing the overfit to the few-shot data. In the middle block of Table 5, we show the ablation study of freezing low-level pre-trained layers and meta-learn only the high-level layers (e.g. the 4-th residual block $\Theta 4$ of ResNet-12) by the *FT* operations. It is obvious that they all yield inferior performances than using our *SS*. An additional observation is that *SS** performs consistently better than *FT**.

Accuracy gain by HT meta-batch. HT meta-batch is basically a curriculum learning scheme, and can be generalized to the models with different network architectures. Comparing the whole model, i.e. *SS* $[\Theta; \theta]$ - HT meta-batch - ResNet-12(pre) (in Table 2), with the *SS* $[\Theta; \theta]$ in Table 5, we can observe that HT meta-batch brings an accuracy boost in a range of 0.6% \sim 2.2%. It is interesting to observe that on the more challenging FC100 dataset it brings an average gain of 1.85%, more than twice of the 0.9% on miniImageNet. In addition, Figure 6 shows the validation results of the models obtained on different meta-training iterations. It demonstrates that our HT meta-batch consistently achieves higher performances than the conventional meta-batch [7]. This ability is also shown in the first column for *FT* based methods (pre-trained ResNet-12 based MAML [7]).

Speed of convergence of MTL with HT meta-batch. MAML [7] used 240k tasks to achieve the best performance on the miniImageNet. Impressively, our HT meta-batch methods used only around 8k tasks to converge, see Figure 6 (b)-(d) (note that each iteration contains 2 tasks). This advantage is more obvious for FC100 on which HT meta-batch methods need at most 6k tasks, Figure 6(g)-

(i). We attest this to three reasons. First, our methods start from the pre-trained deep neural networks. Second, *SS* needs to learn only $< \frac{2}{9}$ parameters of the number of *FT* parameters. Third, HT meta-batch is a hard negative mining step and brings accelerations by learning the challenging tasks [35].

6 CONCLUSIONS

In this paper, we show that our novel MTL model trained with HT meta-batch learning curriculum achieves the top performance for tackling few-shot learning problems. The key operations of MTL on pre-trained DNN neurons proved to be highly efficient for adapting learning experience to the unseen task. The superiority was particularly achieved in the extreme 1-shot cases on three challenging benchmarks – miniImageNet, tieredImageNet and FC100. The generalization ability of our method is validated by implementing MTL on the classical supervised few-shot models as well as the state-of-the-art semi-supervised few-shot models. The consistent improvements by MTL prove that large-scale pre-trained deep networks can offer a good “knowledge base” to conduct efficient few-shot learning on. In terms of learning scheme, HT meta-batch showed consistently good performance for the ablative models. On the more challenging FC100 benchmark, it showed to be particularly helpful for boosting convergence speed. This design is independent of any specific model or architecture and can be generalized well whenever the hardness of task is easy to evaluate in online iterations.

ACKNOWLEDGMENTS

This research is part of NExT research which is supported by the National Research Foundation, Prime Minister’s Office, Singapore under its IRC@SG Funding Initiative. It is also partially supported by German Research Foundation (DFG CRC 1223), and National Natural Science Foundation of China (61772359).

REFERENCES

- [1] Q. Sun, Y. Liu, T.-S. Chua, and B. Schiele, “Meta-transfer learning for few-shot learning,” in *CVPR*, 2019. 1, 2, 5, 6, 7, 8
- [2] M. Ren, E. Triantafillou, S. Ravi, J. Snell, K. Swersky, J. B. Tenenbaum, H. Larochelle, and R. S. Zemel, “Meta-learning for semi-supervised few-shot classification,” in *ICLR*, 2018. 1, 2, 6, 7, 8, 9, 10
- [3] L. Yann, B. Yoshua, and H. Geoffrey, “Deep learning,” *Nature*, vol. 521(7553), p. 436, 2015. 1

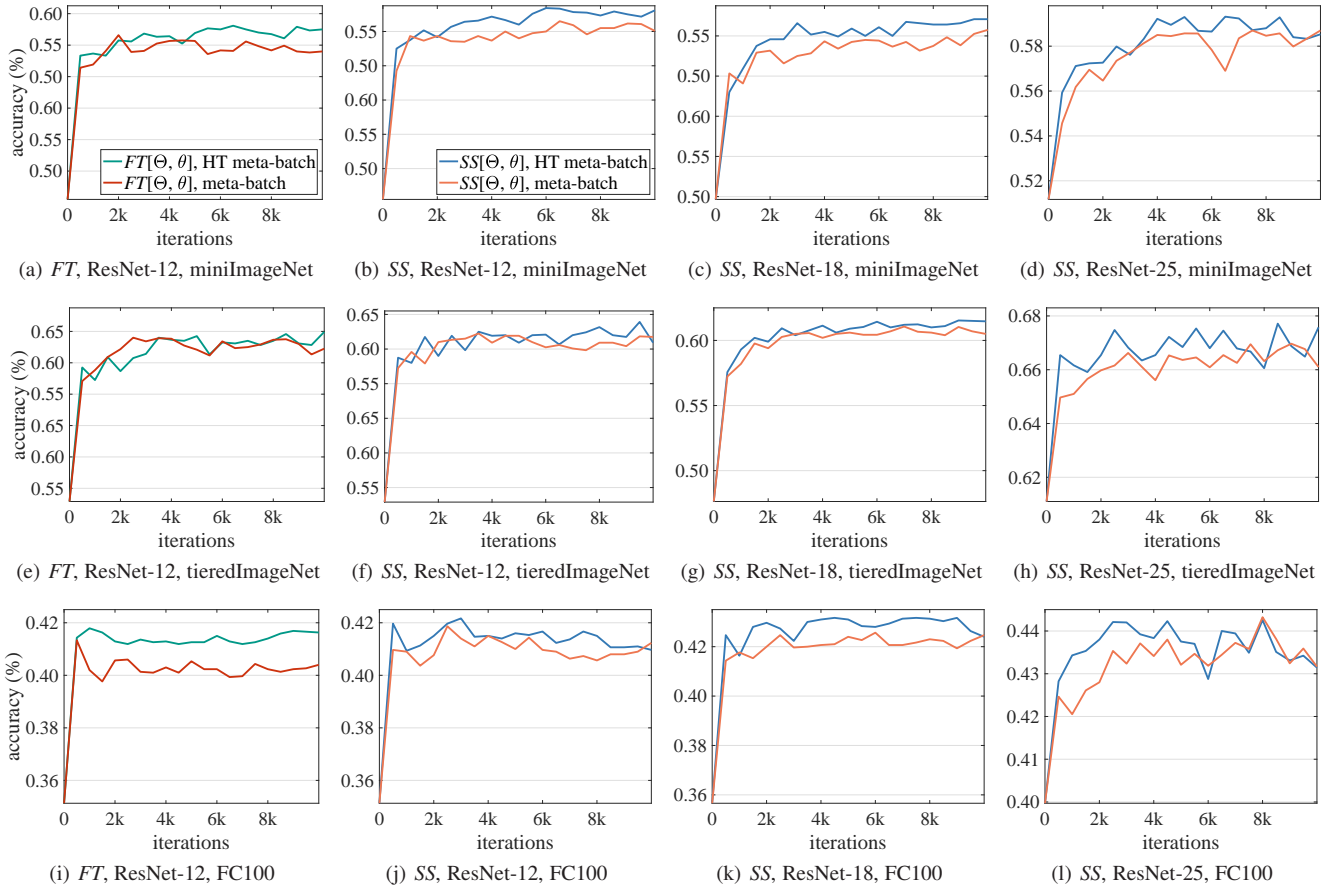


Fig. 6. The 5-way, 1-shot meta-validation accuracy (%) of using FT (deep pre-trained network based MAML [7]) and SS (our MTL) with original meta-batch [7] and our proposed HT meta-batch.

[4] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016. 1, 4, 7

[5] E. Shelhamer, J. Long, and T. Darrell, “Fully convolutional networks for semantic segmentation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 4, pp. 640–651, 2017. 1

[6] F. Li, R. Fergus, and P. Perona, “One-shot learning of object categories,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 4, pp. 594–611, 2006. 1

[7] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *ICML*, 2017. 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12

[8] D. Clevert, T. Unterthiner, and S. Hochreiter, “Fast and accurate deep network learning by exponential linear units (elus),” in *ICLR*, 2016. 1

[9] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang, “Domain adaptation via transfer component analysis,” *IEEE Trans. Neural Networks*, vol. 22, no. 2, pp. 199–210, 2011. 2, 3

[10] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele, “Lucid data dreaming for object tracking,” *arXiv*, vol. 1703.09554, 2017. 1

[11] A. Mehrotra and A. Dukkipati, “Generative adversarial residual pairwise networks for one shot learning,” *arXiv*, vol. 1703.08033, 2017. 1, 2, 9

[12] E. Schwartz, L. Karlinsky, J. Shtok, S. Harary, M. Marder, R. S. Feris, A. Kumar, R. Giryes, and A. M. Bronstein, “Delta-encoder: an effective sample synthesis method for few-shot object recognition,” in *NeurIPS*, 2018. 1, 2, 9

[13] Y. Wang, R. B. Girshick, M. Hebert, and B. Hariharan, “Low-shot learning from imaginary data,” in *CVPR*, 2018. 1, 2

[14] Y. Xian, S. Sharma, B. Schiele, and Z. Akata, “f-VAEGAN-D2: A feature generating framework for any-shot learning,” in *CVPR*, 2019. 1, 2, 3

[15] S. Bartunov and D. P. Vetrov, “Few-shot generative modelling with generative matching networks,” in *AISTATS*, 2018. 1

[16] S. Bengio, Y. Bengio, J. Cloutier, and J. Gecsei, “On the optimization of a synaptic learning rule,” in *Optimality in Artificial and Biological Neural Networks*. Univ. of Texas, 1992, pp. 6–8. 1

[17] D. K. Naik and R. Mammone, “Meta-neural networks that learn by learning,” in *IJCNN*, 1992. 1

[18] S. Thrun and L. Pratt, “Learning to learn: Introduction and overview,” in *Learning to learn*. Springer, 1998, pp. 3–17. 1, 2

[19] C. Finn, K. Xu, and S. Levine, “Probabilistic model-agnostic meta-learning,” in *NeurIPS*, 2018. 1, 5

[20] E. Grant, C. Finn, S. Levine, T. Darrell, and T. L. Griffiths, “Recasting gradient-based meta-learning as hierarchical bayes,” in *ICLR*, 2018. 1, 3, 7, 9

[21] L. Franceschi, P. Frasconi, S. Salzo, R. Grazzi, and M. Pontil, “Bilevel programming for hyperparameter optimization and meta-learning,” in *ICML*, 2018. 1, 2, 3, 7, 9

[22] Y. Lee and S. Choi, “Gradient-based meta-learning with learned layer-wise metric and subspace,” in *ICML*, 2018. 1

[23] R. Zhang, T. Che, Z. Grahahramani, Y. Bengio, and Y. Song, “Metagan: An adversarial approach to few-shot learning,” in *NeurIPS*, 2018. 1, 3, 9

[24] A. Antoniou, H. Edwards, and A. Storkey, “How to train your maml,” in *ICLR*, 2019. 1, 8, 9

[25] A. A. Rusu, D. Rao, J. Sygnowski, O. Vinyals, R. Pascanu, S. Osindero, and R. Hadsell, “Meta-learning with latent embedding optimization,” in *ICLR*, 2019. 1, 3, 6, 9

[26] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, “Matching networks for one shot learning,” in *NIPS*, 2016. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10

[27] T. Munkhdalai and H. Yu, “Meta networks,” in *ICML*, 2017. 2, 9

[28] Z. Li, F. Zhou, F. Chen, and H. Li, “Meta-sgd: Learning to learn quickly for few shot learning,” in *ICML*, 2018. 2

[29] D. Lopez-Paz and M. Ranzato, “Gradient episodic memory for continual learning,” in *NIPS*, 2017. 2, 5

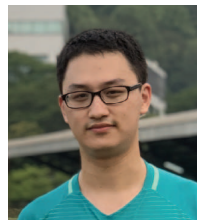
[30] M. McCloskey and N. J. Cohen, “Catastrophic interference in connectionist networks: The sequential learning problem,” in *Psychology of learning and motivation*, 1989, pp. 3–17. 2, 5

[31] J. Snell, K. Swersky, and R. S. Zemel, “Prototypical networks for few-shot learning,” in *NIPS*, 2017. 2, 5, 6, 7, 8, 9, 10

- [32] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. S. Torr, and T. M. Hospedales, "Learning to compare: Relation network for few-shot learning," in *CVPR*, 2018. [2](#), [5](#), [6](#), [8](#), [9](#), [10](#)
- [33] W.-Y. Chen, Y.-C. Liu, Z. Kira, Y.-C. Wang, and J.-B. Huang, "A closer look at few-shot classification," in *ICLR*, 2019. [2](#), [5](#), [6](#), [8](#), [10](#)
- [34] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *ICML*, 2009. [2](#), [3](#)
- [35] A. Shrivastava, A. Gupta, and R. B. Girshick, "Training region-based object detectors with online hard example mining," in *CVPR*, 2016. [2](#), [3](#), [11](#)
- [36] B. N. Oreshkin, P. Rodríguez, and A. Lacoste, "TADAM: task dependent adaptive metric for improved few-shot learning," in *NeurIPS*, 2018. [2](#), [3](#), [7](#), [8](#), [9](#)
- [37] H. E. Geoffrey and P. C. David, "Using fast weights to deblur old memories," in *CogSci*, 1987. [2](#)
- [38] L. Bertinetto, J. F. Henriques, P. H. S. Torr, and A. Vedaldi, "Meta-learning with differentiable closed-form solvers," in *ICLR*, 2019. [2](#), [9](#)
- [39] V. G. Satorras and J. B. Estrach, "Few-shot learning with graph neural networks," in *ICLR*, 2018. [2](#), [9](#)
- [40] Y. Liu, J. Lee, M. Park, S. Kim, and Y. Yang, "Learning to propagate labels: Transductive propagation network for few-shot learning," in *ICLR*, 2019. [2](#), [6](#), [8](#), [9](#), [10](#)
- [41] H. Li, W. Dong, X. Mei, C. Ma, F. Huang, and B. Hu, "Lgm-net: Learning to generate matching networks for few-shot learning," in *ICML*, 2019, pp. 3825–3834. [2](#), [8](#), [9](#)
- [42] H. Li, D. Eigen, S. Dodge, M. Zeiler, and X. Wang, "Finding task-relevant features for few-shot learning by category traversal," in *CVPR*, 2019. [2](#), [6](#), [7](#), [8](#), [9](#)
- [43] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. P. Lillicrap, "Meta-learning with memory-augmented neural networks," in *ICML*, 2016. [2](#)
- [44] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel, "Snail: A simple neural attentive meta-learner," in *ICLR*, 2018. [2](#), [3](#), [7](#), [9](#)
- [45] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in *ICLR*, 2017. [3](#), [7](#), [9](#)
- [46] K. Lee, S. Maji, A. Ravichandran, and S. Soatto, "Meta-learning with differentiable convex optimization," in *CVPR*, 2019. [3](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [47] J. Yang, R. Yan, and A. G. Hauptmann, "Adapting SVM classifiers to data with shifted distributions," in *ICDM Workshops*, 2007. [3](#)
- [48] Y. Wei, Y. Zhang, J. Huang, and Q. Yang, "Transfer learning via learning to transfer," in *ICML*, 2018. [3](#)
- [49] A. R. Zamir, A. Sax, W. B. Shen, L. J. Guibas, J. Malik, and S. Savarese, "Taskonomy: Disentangling task transfer learning," in *CVPR*, 2018. [3](#)
- [50] Q. Sun, B. Schiele, and M. Fritz, "A domain based approach to social relation recognition," in *CVPR*, 2017. [3](#)
- [51] D. Erhan, Y. Bengio, A. C. Courville, P. Manzagol, P. Vincent, and S. Bengio, "Why does unsupervised pre-training help deep learning?" *Journal of Machine Learning Research*, vol. 11, pp. 625–660, 2010. [3](#)
- [52] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, "Speed/accuracy trade-offs for modern convolutional object detectors," in *CVPR*, 2017. [3](#)
- [53] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, "Mask R-CNN," in *ICCV*, 2017. [3](#)
- [54] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, 2018. [3](#)
- [55] M. Rohrbach, S. Ebert, and B. Schiele, "Transfer learning in a transductive setting," in *NIPS*, 2013, pp. 46–54. [3](#)
- [56] R. Keshari, M. Vatsa, R. Singh, and A. Noore, "Learning structure and strength of CNN filters for small sample size training," in *CVPR*, 2018. [3](#)
- [57] S. Qiao, C. Liu, W. Shen, and A. L. Yuille, "Few-shot image recognition by predicting parameters from activations," in *CVPR*, 2018. [3](#), [7](#)
- [58] T. R. Scott, K. Ridgeway, and M. C. Mozer, "Adapted deep embeddings: A synthesis of methods for k-shot inductive transfer learning," in *NeurIPS*, 2018. [3](#)
- [59] N. Sarafianos, T. Giannakopoulos, C. Nikou, and I. A. Kakadiaris, "Curriculum learning for multi-task classification of visual attributes," in *ICCV Workshops*, 2017. [3](#)
- [60] D. Weinshall, G. Cohen, and D. Amir, "Curriculum learning by transfer learning: Theory and experiments with deep networks," in *ICML*, 2018. [3](#)
- [61] A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu, "Automated curriculum learning for neural networks," in *ICML*, 2017. [3](#)
- [62] M. P. Kumar, B. Packer, and D. Koller, "Self-paced learning for latent variable models," in *NIPS*, 2010, pp. 1189–1197. [3](#)
- [63] A. Pentina, V. Sharmanska, and C. H. Lampert, "Curriculum learning of multiple tasks," in *CVPR*, 2015. [3](#)
- [64] L. Jiang, Z. Zhou, T. Leung, L. Li, and L. Fei-Fei, "Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels," in *ICML*, 2018, pp. 2309–2318. [3](#)
- [65] O. Canévet and F. Fleuret, "Large scale hard sample mining with monte carlo tree search," in *CVPR*, 2016. [3](#)
- [66] B. Harwood, V. Kumar, G. Carneiro, I. Reid, and T. Drummond, "Smart mining for deep metric learning," in *ICCV*, 2017. [3](#)
- [67] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *CVPR*, 2005. [3](#)
- [68] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015. [4](#), [7](#)
- [69] S. Gidaris and N. Komodakis, "Dynamic few-shot visual learning without forgetting," in *CVPR*, 2018. [5](#)
- [70] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler, "Rapid adaptation with conditionally shifted neurons," in *ICML*, 2018. [7](#), [9](#)
- [71] A. Krizhevsky, "Learning multiple layers of features from tiny images," *University of Toronto*, 2009. [7](#)
- [72] H. Ye, H. Hu, D. Zhan, and F. Sha, "Learning embedding adaptation for few-shot learning," *arXiv*, vol. 1812.03664, 2018. [7](#)
- [73] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, 2015. [7](#)
- [74] H.-J. Ye, H. Hu, D.-C. Zhan, and F. Sha, "Learning embedding adaptation for few-shot learning," *arXiv preprint arXiv:1812.03664*, 2018. [7](#)
- [75] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv*, vol. 1412.6980, 2014. [7](#), [8](#)



Qianru Sun is an Assistant Professor in the School of Information Systems, Singapore Management University (since Aug 2019). Before that, she was a research fellow working with Prof. Tat-Seng Chua and leads the computer vision group in the NExT++ Lab at the National University of Singapore (since Apr 2018). Before that, she held the Lise Meitner Award Fellowship and was a post-doctoral researcher working with Prof. Dr. Bernt Schiele and Dr. Mario Fritz at the Max-Planck Institute for Informatics (MPII) for two years. She now has a visiting researcher position at MPII. In Jan 2016, she obtained her Ph.D. degree from Peking University, and her thesis was advised by Prof. Hong Liu. From Sep 2014 to Jan 2015, she was a visiting Ph.D. student advised by Prof. Tatsuya Harada at the University of Tokyo. Her research interests are computer vision and machine learning. Specific topics include image recognition, conditional image generation, meta-learning, and transfer learning.



Yaoyao Liu is a Ph.D. student at the School of Electrical and Information Engineering, Tianjin University advised by Prof. Yuting Su and Prof. An-An Liu. From June 2018, he is a visiting Ph.D. student at the School of Computing, National University of Singapore advised by Prof. Tat-Seng Chua and Prof. Qianru Sun. Before this, He obtained his bachelor degree at Qiushi Honors College, Tianjin University. His research lies in few-shot learning, meta learning and conditional image generation.



Zhaozheng Chen is currently an intern and going to be a PhD student (from Jan 2020) at the School of Information Systems, Singapore Management University, advised by Prof. Qianru Sun. Before this, he obtained his bachelor degree at the School of Computer Science and Technology, Shandong University. His research interests are computer vision and machine learning.



Tat-Seng Chua is the KITHCT Chair Professor at the School of Computing, National University of Singapore. He is also the distinguish Visiting Professor of Tsinghua University. He was the Founding Dean of the School from 1998-2000. He is now the Director of a joint research Center between NUS and Tsinghua (NExT) to research into big unstructured multi-source multimodal data analytics. He holds a PhD degree from the University of Leeds, UK, since Feb 1983. Before that, he got his bachelor degree in

the Civil Engineering and Computer Science, University of Leeds, UK, in Jun 1979. His main research interests are in multimedia information retrieval and social media analytics. In particular, his research focuses on the extraction, retrieval and question-answering of text, video and live media arising from the Web and social networks.



Bernt Schiele has been Max Planck Director at MPI for Informatics and Professor at Saarland University since 2010. He studied computer science at the University of Karlsruhe, Germany. He worked on his master thesis in the field of robotics in Grenoble, France, where he also obtained the "diplome d'etudes approfondies d'informatique". In 1994 he worked in the field of multi-modal human-computer interfaces at Carnegie Mellon University, Pittsburgh, PA, USA in the group of Alex Waibel. In 1997 he

obtained his PhD from INP Grenoble, France under the supervision of Prof. James L. Crowley in the field of computer vision. The title of his thesis was "Object Recognition using Multidimensional Receptive Field Histograms". Between 1997 and 2000 he was postdoctoral associate and Visiting Assistant Professor with the group of Prof. Alex Pentland at the Media Laboratory of the Massachusetts Institute of Technology, Cambridge, MA, USA. From 1999 until 2004 he was Assistant Professor at the Swiss Federal Institute of Technology in Zurich (ETH Zurich). Between 2004 and 2010 he was Full Professor at the computer science department of TU Darmstadt.