

Meta-Transfer Learning for Few-Shot Learning

Qianru Sun^{1,3*} Yaoyao Liu^{2*} Tat-Seng Chua¹ Bernt Schiele³

¹National University of Singapore ²Tianjin University[†]

³Max Planck Institute for Informatics, Saarland Informatics Campus

{qsun, schiele}@mpi-inf.mpg.de

liuyaoyao@tju.edu.cn {dcssq, dcscts}@nus.edu.sg

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, meta-learning typically uses shallow neural networks (SNNs), thus limiting its effectiveness. In this paper we propose a novel few-shot learning method called **meta-transfer learning (MTL)** which learns to adapt 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 for MTL. We conduct experiments using (5-class, 1-shot) and (5-class, 5-shot) recognition tasks on two challenging few-shot learning benchmarks: *miniImageNet* and *Fewshot-CIFAR100*. Extensive comparisons to related works validate that our **meta-transfer learning** 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 accuracy¹.

1. Introduction

While deep learning systems have achieved great performance when sufficient amounts of labeled data are available [57, 17, 45], 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 few labeled examples, e.g. 1-shot learn-

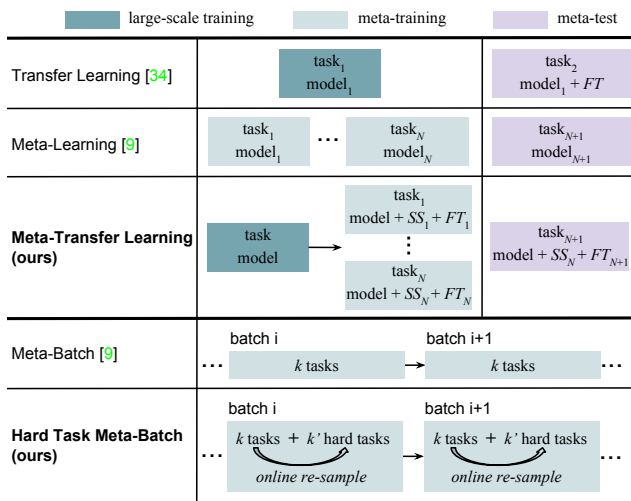


Figure 1. Meta-transfer learning (MTL) is our meta-learning paradigm and hard task (HT) meta-batch is our training strategy. The upper three rows show the differences between MTL and related methods, transfer-learning [34] and meta-learning [9]. The bottom rows compare HT meta-batch with the conventional meta-batch [9]. *FT* stands for fine-tuning a classifier. *SS* represents the *Scaling* and *Shifting* operations in our MTL method.

ing [24]. While humans tend to be highly effective in this context, often grasping the essential connection between new concepts and their own knowledge and experience, it remains challenging for machine learning approaches. E.g., on the CIFAR-100 dataset, a state-of-the-art method [33] achieves only 40.1% accuracy for 1-shot learning, compared to 75.7% for the all-class fully supervised case [6].

Few-shot learning methods can be roughly categorized into two classes: data augmentation and task-based meta-learning. Data augmentation is a classic technique to increase the amount of available data and thus also useful for few-shot learning [21]. Several methods propose to learn a data generator e.g. conditioned on Gaussian noise [28, 43, 53]. However, the generation models often

*Equal contribution.

[†]Yaoyao Liu did this work during his internship at NUS.

¹Code: <https://github.com/y2l/meta-transfer-learning-tensorflow>

underperform when trained on few-shot data [1]. An alternative is to merge data from multiple tasks which, however, is not effective due to variances of the data across tasks [53].

In contrast to data-augmentation methods, meta-learning is a task-level learning method [2, 32, 51]. Meta-learning aims to accumulate experience from learning multiple tasks [9, 38, 47, 30, 13], while base-learning focuses on modeling the data distribution of a single task. A state-of-the-art representative of this, namely Model-Agnostic Meta-Learning (MAML), learns to search for the optimal initialization state to fast adapt a base-learner to a new task [9]. Its task-agnostic property makes it possible to generalize to few-shot supervised learning as well as unsupervised reinforcement learning [13, 10]. 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) to avoid model overfitting, thus being unable to use deeper and more powerful architectures. For example, for the mini-ImageNet dataset [52], MAML uses a *shallow* CNN with only 4 CONV layers and its optimal performance was obtained learning on 240k tasks.

In this paper, we propose a novel meta-learning method called **meta-transfer learning (MTL)** leveraging the advantages of both transfer and meta learning (see conceptual comparison of related methods in Figure 1). In a nutshell, MTL is a novel learning method that helps deep neural nets converge faster while reducing the probability to overfit when using few labeled training 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 [30, 25]. Large-scale trained DNN weights offer a good initialization, enabling fast convergence of meta-transfer learning with fewer tasks, e.g. only 8k tasks for miniImageNet [52], 30 times fewer than MAML [9]. 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. In addition, 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 [26, 27].

The second main contribution of this paper is an effective meta-training curriculum. Curriculum learning [3] and hard negative mining [46] both suggest that faster convergence and stronger performance can be achieved by a 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 [9], 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 two few-shot learning benchmarks, namely miniImageNet [52] and Fewshot-CIFAR100 (FC100) [33], and achieve the state-of-the-art performance.

2. Related work

Few-shot learning Research literature on few-shot learning exhibits great diversity. In this section, we focus on methods using the supervised meta-learning paradigm [12, 51, 9] most relevant to ours and compared to in the experiments. We can divide these methods into three categories. 1) *Metric learning* methods [52, 47, 50] learn a similarity space in which learning is efficient for few-shot examples. 2) *Memory network* methods [30, 41, 33, 29] learn to store “experience” when learning seen tasks and then generalize that to unseen tasks. 3) *Gradient descent* based methods [9, 38, 23, 13, 59] have a specific *meta-learner* that learns to adapt a specific *base-learner* (to few-shot examples) through different tasks. E.g. MAML [9] uses a meta-learner that learns to effectively initialize a base-learner for a new learning task. Meta-learner optimization is done by gradient descent using the validation loss of the base-learner. Our method is closely related. An important difference is that our MTL approach 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 *What* and *how* to transfer are key issues to be addressed in transfer learning, as different methods are applied to different source-target domains and bridge different transfer knowledge [34, 56, 54, 58]. 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 [8]. Another popular transfer method is taking pre-trained networks as backbone and adding high-level functions, e.g. for object detection and recognition [18, 49, 48] and image segmentation [16, 5]. Our meta-transfer learning leverages the idea of transferring pre-trained weights and aims to meta-learn how to effectively transfer. In this paper, large-scale trained DNN weights are *what* to transfer, and the operations of

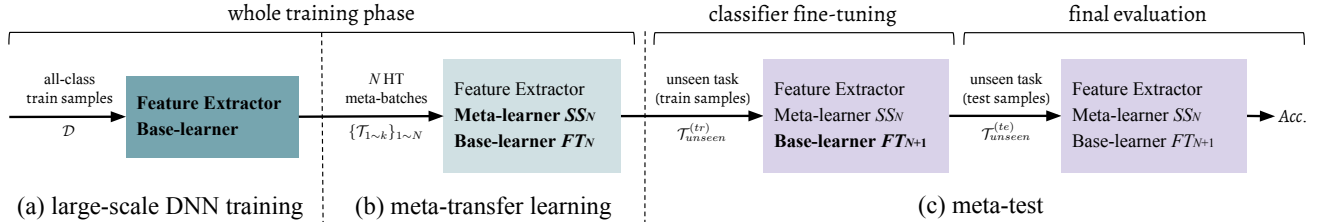


Figure 2. The pipeline of our proposed few-shot learning method, including three phases: (a) DNN training on large-scale data, i.e. using all training datapoints (Section 4.1); (b) Meta-transfer learning (MTL) that learns the parameters of *Scaling* and *Shifting* (SS), based on the pre-trained feature extractor (Section 4.2). Learning is scheduled by the proposed HT meta-batch (Section 4.3); 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. Specifically, SS parameters are learned by meta-training but fixed during meta-test. Base-learner parameters are optimized for every task.

Scaling and *Shifting* indicate *how* to transfer. Similar operations have been used to modulating the per-feature-map distribution of activations for visual reasoning [36].

Some few-shot learning methods have been proposed to use pre-trained weights as initialization [20, 29, 37, 44, 40]. Typically, weights are 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.

Curriculum learning & Hard sample mining Curriculum learning was proposed by Bengio *et al.* [3] and is popular for multi-task learning [35, 42, 55, 14]. 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. Pentina *et al.* [35] use adaptive SVM classifiers to evaluate task difficulty for later organization. Differently, our MTL method does task evaluation online at the phase of episode test, without needing any auxiliary model.

Hard sample mining was proposed by Shrivastava *et al.* [46] for object detection. It treats image proposals overlapped with ground truth as hard negative samples. Training on more confusing data enables the model to achieve higher robustness and better performance [4, 15, 7]. 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 enhance other meta-learning methods, e.g. MAML [9].

3. Preliminary

We introduce the problem setup and notations of meta-learning, following related work [52, 38, 9, 33].

Meta-learning consists of two phases: meta-train and meta-test. A meta-training 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, and a test split $\mathcal{T}^{(te)}$ to optimize the meta-learner. In particular, meta-training 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 $\mathcal{T}_{unseen}^{(te)}$.

Meta-training 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 adapt to the objective of \mathcal{T}_{unseen} by some means, e.g. through initialization [9]. Then, the test result on $\mathcal{T}_{unseen}^{(te)}$ is used to evaluate the meta-learning approach. If there are multiple unseen tasks $\{\mathcal{T}_{unseen}\}$, 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 (64-class, 600-shot) [52], 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 improved overall learning, we use our HT meta-batch strategy (Section 4.3). The training steps are detailed in Algorithm 1 in Section 4.4. Finally, the typical meta-test phase is performed, as introduced in Section 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 [39]. 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 [52], there are totally 64 classes in the training split of \mathcal{D} and each class contains 600 samples used to pre-train a 64-class classifier.

We first randomly initialize a feature extractor Θ (e.g. CONV layers in ResNets [17]) and a classifier θ (e.g. the last FC layer in ResNets [17]), 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 [52].

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 4.3). Figure 3 visualizes the difference of updating through SS and 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 detail the SS operations. 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 of 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,

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

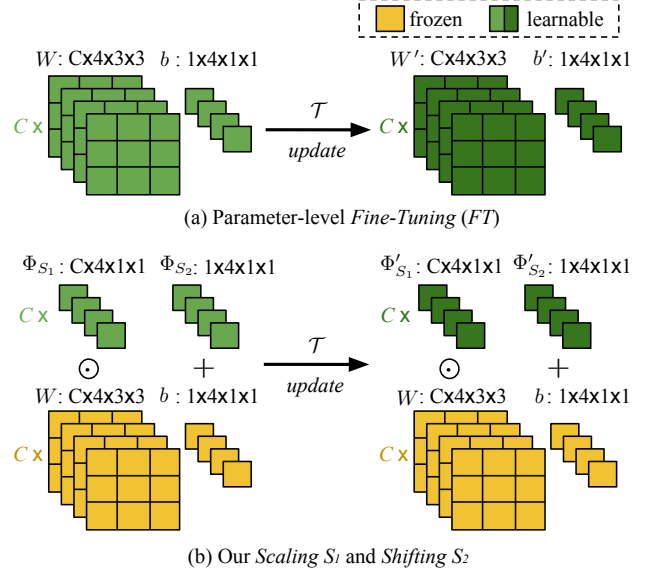


Figure 3. (a) Parameter-level *Fine-Tuning* (FT) is a conventional meta-training operation, e.g. in MAML [9]. Its update works for all neuron parameters, W and b . (b) Our neuron-level *Scaling* and *Shifting* (SS) operations in MTL. They reduce the number of learning parameters and avoid overfitting problems. In addition, they keep large-scale trained parameters (in yellow) frozen, preventing “catastrophic forgetting” [26, 27].

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 3(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 3(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 4.4.

Figure 3(a) shows a typical parameter-level *Fine-Tuning* (FT) operation, which is in the meta optimization phase of our related work MAML [9]. 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’’ [26, 27] when learning specific tasks in MTL. 3) It is light-weight, reducing the chance of overfitting of MTL in few-shot scenarios.

4.3. 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 [9]. 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. We aim to force our meta-learner to ‘‘grow up through hardness’’.

Pipeline. Each task \mathcal{T} has 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. We can also get the recognition accuracy of $\mathcal{T}^{(te)}$ 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 (indexed by $\{m\}$) from k tasks in current meta-batch $\{\mathcal{T}_{1\sim k}\}$, we re-sample tasks from their data. 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 class- m in the current task \mathcal{T} , or (2) indirectly using the label of class- m to sample new samples of that class. In fact, setting (2) considers to include more data variance of class- m and it works better than setting (1) in general.

4.4. Algorithm

Algorithm 1 summarizes the training process of two main stages: large-scale DNN training (line 1-5) and meta-transfer learning (line 6-22). HT meta-batch re-sampling and continuous training phases are shown in lines 16-20, for which the failure classes are returned by Algorithm 2, see line 14. Algorithm 2 presents the learning process on a single task that includes episode training (lines 2-5) and episode test, i.e. meta-level update (lines 6). In lines 7-11,

the recognition rates of all test classes are computed and returned to Algorithm 1 (line 14) for hard task sampling.

Algorithm 1: Meta-transfer learning (MTL)

Input: Task distribution $p(\mathcal{T})$ and corresponding dataset \mathcal{D} , learning rates α, β and γ
Output: Feature extractor Θ , base learner θ , SS 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;
- 7 Reset and re-initialize θ for few-shot tasks;
- 8 **for** meta-batches **do**
- 9 Randomly sample tasks $\{\mathcal{T}\}$ from $p(\mathcal{T})$;
- 10 **while** not done **do**
- 11 Sample task $\mathcal{T}_i \in \{\mathcal{T}\}$;
- 12 Optimize $\Phi_{S_{\{1,2\}}}$ and θ with \mathcal{T}_i by **Algorithm 2**;
- 13 Get the returned class- m then add it to $\{m\}$;
- 14 **end**
- 15 Sample hard tasks $\{\mathcal{T}^{hard}\}$ from $\subseteq p(\mathcal{T}|\{m\})$;
- 16 **while** not done **do**
- 17 Sample task $\mathcal{T}_j^{hard} \in \{\mathcal{T}^{hard}\}$;
- 18 Optimize $\Phi_{S_{\{1,2\}}}$ and θ with \mathcal{T}_j^{hard} by **Algorithm 2**;
- 19 **end**
- 20 Empty $\{m\}$.
- 21 **end**

Algorithm 2: Detail learning steps within a task \mathcal{T}

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

- 1 Sample $\mathcal{T}^{(tr)}$ and $\mathcal{T}^{(te)}$ 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 **while** not done **do**
- 8 Sample class- k in $\mathcal{T}^{(te)}$;
- 9 Compute Acc_k for $\mathcal{T}^{(te)}$;
- 10 **end**
- 11 Return class- m with the lowest accuracy Acc_m .

5. Experiments

We evaluate the proposed **MTL** and **HT meta-batch** in terms of few-shot recognition accuracy and model convergence speed. Below we describe the datasets and detailed settings, followed by an ablation study and a comparison to state-of-the-art methods.

5.1. Datasets and implementation details

We conduct few-shot learning experiments on two benchmarks, miniImageNet [52] and Fewshot-CIFAR100 (FC100) [33]. miniImageNet is widely used in related works [9, 38, 13, 11, 31]. FC100 is newly proposed in [33] and is more challenging in terms of lower image resolution and stricter training-test splits than miniImageNet.

miniImageNet was proposed by Vinyals *et al.* [52] for few-shot learning evaluation. Its complexity is high due to the use of ImageNet images, but requires less resource and infrastructure than running on the full ImageNet dataset [39]. 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 [9, 38, 13, 11, 31].

Fewshot-CIFAR100 (FC100) is based on the popular object classification dataset CIFAR100 [22]. The splits were proposed by [33] (Please check details in the supplementary). It offers a more challenging scenario with lower image resolution and more challenging meta-training/test splits that are separated according to object super-classes. It contains 100 object classes and each class has 600 samples of 32×32 color images. The 100 classes belong 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. These splits accord to super-classes, thus minimize the information overlap between training and val/test tasks.

The following settings are used on both datasets. We train a large-scale DNN with all training datapoints (Section 4.1) and stop this training after $10k$ iterations. We use the same task sampling method as related works [9, 38]. Specifically, 1) we consider the 5-class classification and 2) we sample 5-class, 1-shot (5-shot or 10-shot) episodes to contain 1 (5 or 10) samples for train episode, and 15 (uniform) samples for episode test. Note that in the state-of-the-art work [33], 32 and 64 samples are respectively used in 5-shot and 10-shot settings for episode test. In total, we sample $8k$ tasks for meta-training (same for w/ and w/o HT meta-batch), and respectively sample 600 random tasks for meta-validation and meta-test. Please check the supplementary document (or [GitHub repository](#)) for other implementation details, e.g. learning rate and dropout rate.

Network architecture. We present the details for the Fea-

ture Extractor Θ , MTL meta-learner with *Scaling* Φ_{S_1} and *Shifting* Φ_{S_2} , and MTL base-learner (classifier) θ .

The architecture of Θ have two options, ResNet-12 and 4CONV, commonly used in related works [9, 52, 38, 31, 29, 33]. **4CONV** consists of 4 layers with 3×3 convolutions and 32 filters, followed by batch normalization (BN) [19], a ReLU nonlinearity, and 2×2 max-pooling. **ResNet-12** is more popular in recent works [33, 29, 11, 31]. It 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. **The difference** between using 4CONV and using ResNet-12 in our methods is that ResNet-12 MTL sees the large-scale data training, but 4CONV MTL is learned from scratch because of its poor performance for large-scale data training (see results in the supplementary). Therefore, we emphasize the experiments of using ResNet-12 MTL for its superior performance. **The architectures of Φ_{S_1} and Φ_{S_2}** are generated according to the architecture of Θ , as introduced in Section 4.2. That is when using ResNet-12 in MTL, Φ_{S_1} and Φ_{S_2} also have 12 layers, respectively. **The architecture of θ** is an FC layer. We empirically find that a single FC layer is faster to train and more effective for classification than multiple layers. (see comparisons in the supplementary).

5.2. Ablation study setting

In order to show the effectiveness of our approach, we design some ablative settings: two baselines without meta-learning but more classic learning, three baselines of *Fine-Tuning (FT)* on smaller number of parameters (Table 1), and two MAML variants using our deeper pre-trained model and HT meta-batch (Table 2 and Table 3). Note that the alternative meta-learning operation to *SS* is the *FT* used in MAML. Some bullet names 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 $[\Theta_4; \theta]$ (or θ). These are straight-forward ways to define a smaller set of meta-learner parameters than MAML. We can freeze low-level pre-trained layers and meta-learn the classifier layer θ with (or without) high-level CONV layer Θ_4 that is the 4th residual block of ResNet-12.

5.3. Results and analysis

Table 1, Table 2 and Table 3 present the overall results on miniImageNet and FC100 datasets. Extensive comparisons are done with ablative methods and the state-of-the-arts. Note that tables present the highest accuracies for which the iterations were chosen by validation. For the miniIma-

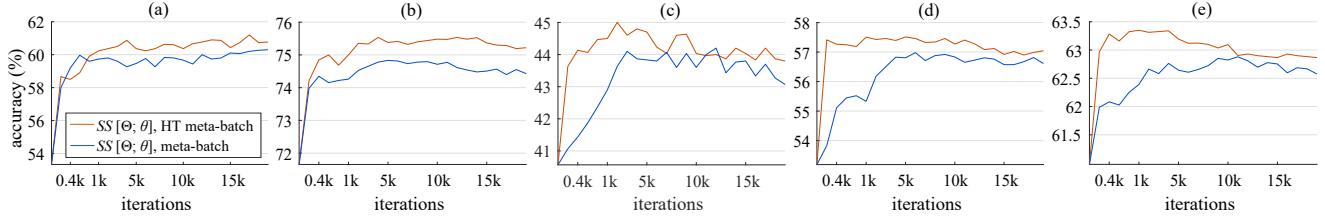


Figure 4. (a)(b) show the results of 1-shot and 5-shot on miniImageNet; (c)(d)(e) show the results of 1-shot, 5-shot and 10-shot on FC100.

geNet, iterations for 1-shot and 5-shot are at 17k and 14k, respectively. For the FC100, iterations are all at 1k. Figure 4 shows the performance gap between *with* and *without* HT meta-batch in terms of accuracy and converging speed.

Result overview on miniImageNet. In Table 2, we can see that the proposed MTL with $SS [\Theta; \theta]$, HT meta-batch and ResNet-12(pre) achieves the best few-shot classification performance with 61.2% for (5-class, 1-shot). Besides, it tackles the (5-class, 5-shot) tasks with an accuracy of 75.5% that is comparable to the state-of-the-art results, i.e. 76.7%, reported by TADAM [33] whose model used 72 additional FC layers in the ResNet-12 arch. In terms of the network arch, it is obvious that models using ResNet-12 (pre) outperforms those using 4CONV by large margins, e.g. 4CONV models have the best 1-shot result with 50.44% [50] which is 10.8% lower than our best.

Result overview on FC100. In Table 3, we give the results of TADAM using their reported numbers in the paper [33]. We used the public code of MAML [9] to get its results for this new dataset. Comparing these methods, we can see that MTL consistently outperforms MAML by large margins, i.e. around 7% in all tasks; and surpasses TADAM by a relatively larger number of 5% for 1-shot, and with 1.5% and 1.8% respectively for 5-shot and 10-shot tasks.

MTL vs. No meta-learning. Table 1 shows the results of *No meta-learning* on the top block. Compared to these, our approach achieves significantly better performance even *without* HT meta-batch, e.g. the largest margins are 10.2% for 1-shot and 8.6% for 5-shot on miniImageNet. This validates the effectiveness of our meta-learning method for tackling few-shot learning problems. Between two *No*

	miniImageNet		FC100		
	1 (shot)	5	1	5	10
<i>update</i> $[\Theta; \theta]$	45.3	64.6	38.4	52.6	58.6
<i>update</i> θ	50.0	66.7	39.3	51.8	61.0
<i>FT</i> θ	55.9	71.4	41.6	54.9	61.1
<i>FT</i> $[\Theta 4; \theta]$	57.2	71.6	40.9	54.3	61.3
<i>FT</i> $[\Theta; \theta]$	58.3	71.6	41.6	54.4	61.2
$SS [\Theta 4; \theta]$	59.2	73.1	42.4	55.1	61.6
$SS [\Theta; \theta]$ (Ours)	60.2	74.3	43.6	55.4	62.4

Table 1. Classification accuracy (%) using ablative models, on two datasets. “meta-batch” and “ResNet-12(pre)” are used.

meta-learning methods, we can see that updating both feature extractor Θ and classifier θ is inferior to updating θ only, 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.

Performance effects of MTL components. MTL with full components, $SS [\Theta; \theta]$, HT meta-batch and ResNet-12(pre), achieves the best performances for all few-shot settings on both datasets, see Table 2 and Table 3. We can conclude that our large-scale network training on deep CNN significantly boost the few-shot learning performance. This is an important gain brought by the transfer learning idea in our MTL approach. It is interesting to note that this gain on FC100 is not as large as for miniImageNet: only 1.7%, 1.0% and 4.0%. The possible reason is that FC100 tasks for meta-train and meta-test are clearly split according to super-classes. The data domain gap is larger than that for miniImageNet, which makes transfer more difficult.

HT meta-batch and ResNet-12(pre) in our approach can be generalized to other meta-learning models. MAML 4CONV with HT meta-batch gains averagely 1% on two datasets. When changing 4CONV by deep ResNet-12 (pre) it achieves significant improvements, e.g. 10% and 9% on miniImageNet. Compared to MAML variants, our MTL results are consistently higher, e.g. 2.5% ~ 3.3% on FC100. People may argue that MAML fine-tuning(*FT*) all network parameters is likely to overfit to few-shot data. In the middle block of Table 1, 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 of ResNet-12) by the *FT* operations of MAML. These all yield inferior performances than using our SS . An additional observation is that SS^* performs consistently better than FT^* .

Speed of convergence of MTL. MAML [9] used 240k tasks to achieve the best performance on miniImageNet. Impressively, our MTL methods used only 8k tasks, see Figure 4(a)(b) (note that each iteration contains 2 tasks). This advantage is more obvious for FC100 on which MTL methods need at most 2k tasks, Figure 4(c)(d)(e). We attest this to two reasons. First, MTL starts from the pre-trained ResNet-12. And second, SS (in MTL) needs to learn only $< \frac{2}{9}$ parameters of the number of *FT* (in MAML) when using ResNet-12.

Few-shot learning method		Feature extractor	1-shot	5-shot
<i>Data augmentation</i>	Adv. ResNet, [28]	WRN-40 (pre)	55.2	69.6
	Delta-encoder, [43]	VGG-16 (pre)	58.7	73.6
<i>Metric learning</i>	Matching Nets, [52]	4 CONV	43.44 ± 0.77	55.31 ± 0.73
	ProtoNets, [47]	4 CONV	49.42 ± 0.78	68.20 ± 0.66
	CompareNets, [50]	4 CONV	50.44 ± 0.82	65.32 ± 0.70
<i>Memory network</i>	Meta Networks, [30]	5 CONV	49.21 ± 0.96	–
	SNAIL, [29]	ResNet-12 (pre) [◊]	55.71 ± 0.99	68.88 ± 0.92
	TADAM, [33]	ResNet-12 (pre) [†]	58.5 ± 0.3	76.7 ± 0.3
<i>Gradient descent</i>	MAML, [9]	4 CONV	48.70 ± 1.75	63.11 ± 0.92
	Meta-LSTM, [38]	4 CONV	43.56 ± 0.84	60.60 ± 0.71
	Hierarchical Bayes, [13]	4 CONV	49.40 ± 1.83	–
	Bilevel Programming, [11]	ResNet-12 [◊]	50.54 ± 0.85	64.53 ± 0.68
	MetaGAN, [59]	ResNet-12	52.71 ± 0.64	68.63 ± 0.67
	adaResNet, [31]	ResNet-12 [‡]	56.88 ± 0.62	71.94 ± 0.57
MAML, HT	$FT [\Theta; \theta]$, HT meta-batch	4 CONV	49.1 ± 1.9	64.1 ± 0.9
MAML deep, HT	$FT [\Theta; \theta]$, HT meta-batch	ResNet-12 (pre)	59.1 ± 1.9	73.1 ± 0.9
MTL (Ours)	$SS [\Theta; \theta]$, meta-batch	ResNet-12 (pre)	60.2 ± 1.8	74.3 ± 0.9
	$SS [\Theta; \theta]$, HT meta-batch	ResNet-12 (pre)	61.2 ± 1.8	75.5 ± 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 dataset. “pre” means pre-trained for a single classification task using all training datapoints.

Few-shot learning method		Feature extractor	1-shot	5-shot	10-shot
<i>Gradient descent</i>	MAML, [9] [‡]	4 CONV	38.1 ± 1.7	50.4 ± 1.0	56.2 ± 0.8
<i>Memory network</i>	TADAM, [33]	ResNet-12 (pre) [†]	40.1 ± 0.4	56.1 ± 0.4	61.6 ± 0.5
MAML, HT	$FT [\Theta; \theta]$, HT meta-batch	4 CONV	39.9 ± 1.8	51.7 ± 0.9	57.2 ± 0.8
MAML deep, HT	$FT [\Theta; \theta]$, HT meta-batch	ResNet-12 (pre)	41.8 ± 1.9	55.1 ± 0.9	61.9 ± 0.8
MTL (Ours)	$SS [\Theta; \theta]$, meta-batch	ResNet-12 (pre)	43.6 ± 1.8	55.4 ± 0.9	62.4 ± 0.8
	$SS [\Theta; \theta]$, HT meta-batch	ResNet-12 (pre)	45.1 ± 1.8	57.6 ± 0.9	63.4 ± 0.8

†Additional 72 fully connected layers ‡Our implementation using the public code of MAML.

Table 3. The 5-way with 1-shot, 5-shot and 10-shot classification accuracy (%) on Fewshot-CIFAR100 (FC100) dataset. “pre” means pre-trained for a single classification task using all training datapoints.

Speed of convergence of HT meta-batch. Figure 4 shows 1) MTL with HT meta-batch consistently achieves higher performances than MTL with the conventional meta-batch [9], in terms of the recognition accuracy in all settings; and 2) it is impressive that MTL with HT meta-batch achieves top performances early, after *e.g.* about $2k$ iterations for 1-shot, $1k$ for 5-shot and $1k$ for 10-shot, on the more challenging dataset – FC100.

6. Conclusions

In this paper, we show that our novel MTL 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 highly efficient for adapting learning experience to the unseen task. The superiority was particularly achieved in the extreme 1-

shot cases on two challenging benchmarks – miniImageNet and FC100. In terms of learning scheme, HT meta-batch showed consistently good performance for all baselines and ablative models. On the more challenging FC100 benchmark, it showed to be particularly helpful for boosting convergence speed. This design is independent from any specific model and could 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] S. Bartunov and D. P. Vetrov. Few-shot generative modelling with generative matching networks. In *AISTATS*, 2018. 2
- [2] 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*, pages 6–8. Univ. of Texas, 1992. 2
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009. 2, 3
- [4] O. Canévet and F. Fleuret. Large scale hard sample mining with monte carlo tree search. In *CVPR*, 2016. 3
- [5] 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.*, 40(4):834–848, 2018. 2
- [6] D. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*, 2016. 1
- [7] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, 2005. 3
- [8] 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*, 11:625–660, 2010. 2
- [9] 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, 6, 7, 8
- [10] C. Finn, K. Xu, and S. Levine. Probabilistic model-agnostic meta-learning. In *NeurIPS*, 2018. 2
- [11] L. Franceschi, P. Frasconi, S. Salzo, R. Grazi, and M. Pontil. Bilevel programming for hyperparameter optimization and meta-learning. In *ICML*, 2018. 6, 8
- [12] H. E. Geoffrey and P. C. David. Using fast weights to deblur old memories. In *CogSci*, 1987. 2
- [13] E. Grant, C. Finn, S. Levine, T. Darrell, and T. L. Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *ICLR*, 2018. 2, 6, 8
- [14] A. Graves, M. G. Bellemare, J. Menick, R. Munos, and K. Kavukcuoglu. Automated curriculum learning for neural networks. In *ICML*, 2017. 3
- [15] B. Harwood, V. Kumar, G. Carneiro, I. Reid, and T. Drummond. Smart mining for deep metric learning. In *ICCV*, 2017. 3
- [16] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick. Mask R-CNN. In *ICCV*, 2017. 2
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 1, 4
- [18] 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. 2
- [19] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. 6
- [20] 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
- [21] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid data dreaming for object tracking. *arXiv*, 1703.09554, 2017. 1
- [22] A. Krizhevsky. Learning multiple layers of features from tiny images. *University of Toronto*, 2009. 6
- [23] Y. Lee and S. Choi. Gradient-based meta-learning with learned layerwise metric and subspace. In *ICML*, 2018. 2
- [24] F. Li, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(4):594–611, 2006. 1
- [25] Z. Li, F. Zhou, F. Chen, and H. Li. Meta-sgd: Learning to learn quickly for few shot learning. In *ICML*, 2018. 2
- [26] D. Lopez-Paz and M. Ranzato. Gradient episodic memory for continual learning. In *NIPS*, 2017. 2, 4, 5
- [27] M. McCloskey and N. J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, pages 3–17, 1989. 2, 4, 5
- [28] A. Mehrotra and A. Dukkupati. Generative adversarial residual pairwise networks for one shot learning. *arXiv*, 1703.08033, 2017. 1, 8
- [29] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. Snail: A simple neural attentive meta-learner. In *ICLR*, 2018. 2, 3, 6, 8
- [30] T. Munkhdalai and H. Yu. Meta networks. In *ICML*, 2017. 2, 8
- [31] T. Munkhdalai, X. Yuan, S. Mehri, and A. Trischler. Rapid adaptation with conditionally shifted neurons. In *ICML*, 2018. 6, 8
- [32] D. K. Naik and R. Mammone. Meta-neural networks that learn by learning. In *IJCNN*, 1992. 2
- [33] B. N. Oreshkin, P. Rodríguez, and A. Lacoste. TADAM: task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018. 1, 2, 3, 6, 7, 8
- [34] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain adaptation via transfer component analysis. *IEEE Trans. Neural Networks*, 22(2):199–210, 2011. 1, 2
- [35] A. Pentina, V. Sharmanska, and C. H. Lampert. Curriculum learning of multiple tasks. In *CVPR*, 2015. 3
- [36] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. C. Courville. Film: Visual reasoning with a general conditioning layer. In *AAAI*, 2018. 3
- [37] S. Qiao, C. Liu, W. Shen, and A. L. Yuille. Few-shot image recognition by predicting parameters from activations. In *CVPR*, 2018. 3
- [38] S. Ravi and H. Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 2, 3, 6, 8
- [39] 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*, 115(3):211–252, 2015. 4, 6
- [40] 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. 3

- [41] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. P. Lillicrap. Meta-learning with memory-augmented neural networks. In *ICML*, 2016. 2
- [42] N. Sarafianos, T. Giannakopoulos, C. Nikou, and I. A. Kakadiaris. Curriculum learning for multi-task classification of visual attributes. In *ICCV Workshops*, 2017. 3
- [43] 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, 8
- [44] 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
- [45] E. Shelhamer, J. Long, and T. Darrell. Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):640–651, 2017. 1
- [46] A. Shrivastava, A. Gupta, and R. B. Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016. 2, 3
- [47] J. Snell, K. Swersky, and R. S. Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. 2, 8
- [48] Q. Sun, L. Ma, S. Joon Oh, L. Van Gool, B. Schiele, and M. Fritz. Natural and effective obfuscation by head inpainting. In *CVPR*, 2018. 2
- [49] Q. Sun, B. Schiele, and M. Fritz. A domain based approach to social relation recognition. In *CVPR*, 2017. 2
- [50] 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, 7, 8
- [51] S. Thrun and L. Pratt. Learning to learn: Introduction and overview. In *Learning to learn*, pages 3–17. Springer, 1998. 2
- [52] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. 2, 3, 4, 6, 8
- [53] Y. Wang, R. B. Girshick, M. Hebert, and B. Hariharan. Low-shot learning from imaginary data. In *CVPR*, 2018. 1, 2
- [54] Y. Wei, Y. Zhang, J. Huang, and Q. Yang. Transfer learning via learning to transfer. In *ICML*, 2018. 2
- [55] D. Weinshall, G. Cohen, and D. Amir. Curriculum learning by transfer learning: Theory and experiments with deep networks. In *ICML*, 2018. 3
- [56] J. Yang, R. Yan, and A. G. Hauptmann. Adapting SVM classifiers to data with shifted distributions. In *ICDM Workshops*, 2007. 2
- [57] L. Yann, B. Yoshua, and H. Geoffrey. Deep learning. *Nature*, 521(7553):436, 2015. 1
- [58] A. R. Zamir, A. Sax, W. B. Shen, L. J. Guibas, J. Malik, and S. Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. 2
- [59] R. Zhang, T. Che, Z. Grahahramani, Y. Bengio, and Y. Song. Metagan: An adversarial approach to few-shot learning. In *NeurIPS*, 2018. 2, 8