

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and  
Information Systems

School of Computing and Information Systems

---

8-2020

### Adaptive task sampling for meta-learning

Chenghao LIU

*Singapore Management University, chliu@smu.edu.sg*

Zhihao WANG

*Singapore Management University, zhwang@smu.edu.sg*

Doyen SAHOO

*Singapore Management University, doyens@smu.edu.sg*

Yuan FANG

*Singapore Management University, yfang@smu.edu.sg*

Kun ZHANG

*See next page for additional authors*

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Databases and Information Systems Commons](#), and the [Theory and Algorithms Commons](#)

---

#### Citation

LIU, Chenghao; WANG, Zhihao; SAHOO, Doyen; FANG, Yuan; ZHANG, Kun; and HOI, Steven C. H.. Adaptive task sampling for meta-learning. (2020). *Proceedings of the 16th European Conference on Computer Vision, Glasgow, UK, 2020, August 23 - 28*. 1-17.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/5293](https://ink.library.smu.edu.sg/sis_research/5293)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

---

**Author**

Chenghao LIU, Zhihao WANG, Doyen SAHOO, Yuan FANG, Kun ZHANG, and Steven C. H. HOI

# Adaptive Task Sampling for Meta-Learning <sup>\*</sup>

Chenghao Liu<sup>1</sup>   Zhihao Wang<sup>2</sup>   Doyen Sahoo<sup>3</sup>   Yuan Fang<sup>1</sup>   Kun Zhang<sup>4</sup>  
Steven C.H. Hoi<sup>1,3</sup>

Singapore Management University<sup>1</sup>   South China University of Technology<sup>2</sup>  
Salesforce Research Asia<sup>3</sup>   Carnegie Mellon University<sup>4</sup>  
{chliu, yfang}@smu.edu.sg, ptkin@outlook.com,  
{dsahoo, shoi}@salesforce.com, kunz1@cmu.edu

**Abstract.** Meta-learning methods have been extensively studied and applied in computer vision, especially for few-shot classification tasks. The key idea of meta-learning for few-shot classification is to mimic the few-shot situations faced at test time by randomly sampling classes in meta-training data to construct few-shot tasks for episodic training. While a rich line of work focuses solely on how to extract meta-knowledge across tasks, we exploit the complementary problem on how to generate informative tasks. We argue that the randomly sampled tasks could be sub-optimal and uninformative (e.g., the task of classifying “dog” from “laptop” is often trivial) to the meta-learner. In this paper, we propose an adaptive task sampling method to improve the generalization performance. Unlike instance based sampling, task based sampling is much more challenging due to the implicit definition of the task in each episode. Therefore, we accordingly propose a greedy class-pair based sampling method, which selects difficult tasks according to class-pair potentials. We evaluate our adaptive task sampling method on two few-shot classification benchmarks, and it achieves consistent improvements across different feature backbones, meta-learning algorithms and datasets.

## 1 Introduction

Deep neural networks have achieved great performance in areas such as image recognition [17], machine translation [9] and speech synthesis [51] when large amounts of labelled data are available. In stark contrast, human intelligence naturally possesses the ability to leverage prior knowledge and quickly learn new concepts from only a handful of samples. Such fast adaptation is made possible by some fundamental structures in human brains such as the “shape bias” to learn the learning procedure [25], which is also known as *meta-learning*. The fact that deep neural networks fail in the small data regime formulates a desirable problem for understanding intelligence. In particular, leveraging meta-learning algorithms to solve few-shot learning problems [24, 38] has recently gained much attention, which aims to close the gap between human and machine intelligence by training deep neural networks that can generalize well from very few labelled samples. In this setup, meta-learning is formulated as the extraction

---

<sup>\*</sup> The first two authors contributed equally, and completed most of this work when working at the School of Information Systems, Singapore Management University (SMU). Steven C.H. Hoi is currently with Salesforce Research Asia and on leave from SMU.

of cross-task knowledge that can facilitate the quick acquisition of task-specific knowledge from new tasks.

In order to compensate for the scarcity of training data in few-shot classification tasks, meta-learning approaches rely on an *episodic training* paradigm. A series of few-shot tasks are sampled from meta-training data for the extraction of transferable knowledge across tasks, which is then applied to new few-shot classification tasks consisting of unseen classes during the meta-testing phase. Specifically, optimization-based meta-learning approaches [46, 12] aim to find a global set of model parameters that can be quickly and effectively fine-tuned for each individual task with just a few gradient descent update steps. Meanwhile, metric-based meta-learning approaches [47, 37] learn a shared distance metric across tasks.

Despite their noticeable improvements, these meta-learning approaches leverage uniform sampling over classes to generate few-shot tasks, which ignores the intrinsic relationships between classes when forming episodes. We argue that exploiting class structures to construct more informative tasks is critical in meta-learning, which improves its ability to adapt to novel classes. For example, in the midst of the training procedure, a randomly sampled task of classifying dogs from laptops may have little effect on the model update due to its simpleness. Furthermore, in the conventional classification problem, prioritizing challenging training examples [43, 42] to improve the generalization performance has been widely used in various fields, ranging from AdaBoost [14] that selects harder examples to train subsequent classifiers, to Focal Loss [28] that adds a soft weighting scheme to emphasize harder examples.

A natural question thus arises: Can we perform adaptive task sampling and create more difficult tasks for meta-learning? Compared to the traditional instance-based adaptive sampling scheme, one key challenge in task sampling is to define the difficulty of a task. A naïve solution is to choose the difficult classes since each task is constructed by multiple classes. However, the difficulty of a class, and even the semantics of a class, is dependent on each other. For instance, the characteristics to discriminate “dog” from “laptop” or “car” are relatively easier to uncover than those for discriminating “dog” from “cat” or “tiger”. In other words, the difficulty of a task goes beyond the difficulty of individual classes, and adaptive task sampling should consider the intricate relationships between different classes.

In this work, we propose a **class-pair based** adaptive task sampling method for meta-learning with several appealing qualities. First, it determines the task selection distribution by computing the difficulty of all class-pairs in it. As a result, it could capture the complex-structured relationships between classes in a multi-class few-shot classification problem. Second, since the cost of computing the task selection distribution for  $K$ -way classification problem is  $(|\mathcal{C}_{tr}| \text{ choose } K)$  or  $O(|\mathcal{C}_{tr}|^K)$ , where  $|\mathcal{C}_{tr}|$  is the number of classes in the meta-training data, we further propose a **greedy class-pair based** adaptive task sampling method which only requires  $O(K)$  time. Meanwhile, it can be formally established that the proposed greedy approach in fact samples from a distribution that is identical to that in the non-greedy version. Lastly, our method could be applied to any meta-learning algorithms that follow episodic training and works well with different feature backbones.

In summary, our work makes the following contributions. **(1)** We propose a class-pair based adaptive task sampling approach for meta-learning methods, to improve the generalization performance on unseen tasks. **(2)** We further develop a greedy class-pair based approach that not only significantly reduces the complexity of task distribution computation, but also guarantees the generation of an identical distribution as that in the non-greedy approach. **(3)** We study the impact of the adaptive task sampling method by integrating it with various meta-learning approaches and performing comprehensive experiments on the miniImageNet and CIFAR-FS few-shot datasets, which quantitatively demonstrates the superior performance of our method. **(4)** We also conduct an extensive investigation of different sampling strategies, including class-based method, easy class-pair based method and uncertain class-pair based method. The results show that hard class-pair based sampling consistently leads to more accurate results.

## 2 Related Work

**Meta-learning:** The original idea of meta-learning, training a meta-model to learn a base model, has existed for at least 20 years [48, 35]. Recently, the meta-learning framework has been used to solve few-shot classification problems. One typical work is the optimization based method. [38] uses the LSTM-based meta-learner to replace the SGD optimizer in the base model. MAML [12] and its variants [27, 4] aim to learn a good model initialization so that the model for new tasks can be learned with a small number of samples and gradient update steps. Another category of work is the metric based method. It learns a set of embedding functions such that when represented in this space, images are easy to be recognized using a non-parametric model like nearest neighbor [50, 44, 37]. All of these methods follow the uniform sampling scheme to generate tasks at each episode. Besides, [46] considers a heuristic sampling method, which uses memory to store all the failure classes from  $k$  continuous tasks, and then constructs a hard task from them. [49, 29] utilize pre-defined class structure information to construct tasks in both meta-training and meta-testing phases. In this way, the experiment setting could more closely resemble realistic scenarios. In contrast, our work, inspired by importance sampling in stochastic optimization, aims to adaptively update task generating distribution in the meta-training phase, and this, in turn, improves its ability to adapt to novel classes with few training data in the meta-testing phase. We also present a theoretical analysis of the generalization bound to justify our approach.

**Adaptive Sampling:** Instance-based sampling is ubiquitous in stochastic optimization. Generally, it constantly reevaluates the relative importance of each instance during training. The most common paradigm is to calculate the importance of each instance based on the gradient norm [1], bound on the gradient norm [20], loss [31], approximate loss [21] or prediction probability [7]. One typical line of research work is to leverage adaptive sampling for fast convergence [55, 2]. Researchers also consider improving the generalization performance rather than speeding up training [30]. Specifically, [5] considers instances that increase difficulty. Hard example mining methods also prioritize challenging training examples [43, 28]. Some other researchers prioritize uncertain examples that are close to the model’s decision boundary [7, 45]. In this

work, we also evaluate easy sampling and uncertain sampling at the task level, but experimental results show that hard sampling performs better. There also exists work for sampling mini-batches instead of a single instance [11, 18]. [52, 53] consider sampling diverse mini-batches via the repulsive point process. Nonetheless, these methods are not designed for meta-learning and few-shot learning.

### 3 Preliminaries

In this section, we review the *episodic training* paradigm in meta-learning and the vanilla instance-based adaptive sampling method for SGD.

#### 3.1 Episodic Training

In the meta-learning problem setting, the goal is to learn models that can learn new tasks from small amounts of data. Formally, we have a large meta-training dataset  $\mathbb{D}_{tr}$  (typically containing a large number of classes) and a meta-test dataset  $\mathbb{D}_{test}$ , in which their respective category sets  $\mathbb{C}_{tr} = \{1, \dots, |\mathbb{C}_{tr}|\}$  and  $\mathbb{C}_{test} = \{|\mathbb{C}_{tr}| + 1, \dots, |\mathbb{C}_{tr}| + |\mathbb{C}_{test}|\}$  are disjoint. We aim to learn a classification model on  $\mathbb{D}_{tr}$  that can generalize to unseen categories  $\mathbb{C}_{test}$  with one or few training examples per category.

The success of existing meta-learning approaches relies on the *episodic training* paradigm [50], which mimics the few-shot regime faced at test time during training on  $\mathbb{D}_{tr}$ . Particularly, meta-learning algorithms learn from a collection of  $K$ -way- $M$ -shot classification tasks sampled from the amply labelled set  $\mathbb{D}_{tr}$  and are evaluated in a similar way on  $\mathbb{D}_{test}$ . In each episode of meta-training, we first sample  $K$  classes  $\mathbb{L}^K \sim \mathbb{C}_{tr}$ . Then, we sample  $M$  and  $N$  labelled images per class in  $\mathbb{L}^K$  to construct the support set  $\mathbb{S} = \{(s_m, y_m)_m\}$  and query set  $\mathbb{Q} = \{(q_n, y_n)_n\}$ , respectively. The episodic training for few-shot learning is achieved by minimizing, for each episode, the loss of the prediction for each sample in the query set, given the support set. The model is parameterized by  $\theta$  and the loss is the negative loglikelihood of the true class of each query sample:

$$\ell(\theta) = \mathbb{E}_{(S, Q)} \left[ - \sum_{(q_n, y_n) \in Q} \log p_\theta(y_n | q_n, S) \right], \quad (1)$$

where  $p_\theta(y_n | q_n, S)$  is the classification probability based on the support set. The model then back-propagates the gradient of the total loss  $\nabla \ell(\theta)$ . Different meta-learning approaches differ in the manner in which this conditioning on the support set is realized. To better explain how it works, we show its framework in Figure 1.

#### 3.2 Instance-base Adaptive Sampling for SGD

Let  $\mathbb{D} = \{(x_i, y_i)_i\}$  indicate the training dataset. The probability of selecting each sample is equal at the initial stage (i.e.,  $p_0(i|\mathbb{D}) = \frac{1}{|\mathbb{D}|}$ ). To emphasize difficult examples while applying SGD, we adaptively update the selection probability  $p^{t+1}(i)$  for

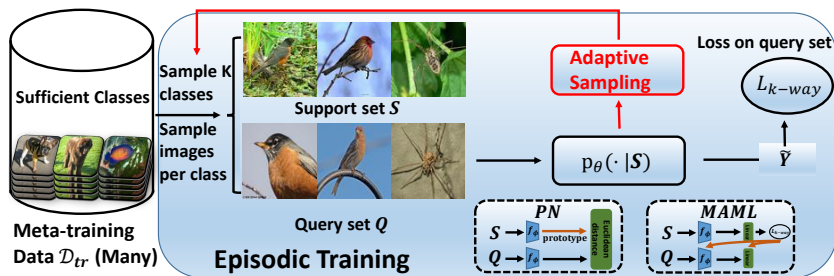


Fig. 1: The episodic training paradigm for meta-learning few-shot classification.

instance  $i$  at iteration  $t + 1$  according to the current prediction probability  $p(y_i|x_i)$  and the selection probability at previous iteration  $p^t(i)$ ,

$$p^{t+1}(i) \propto (p^t(i))^\tau e^{\alpha(1-p(y_i|x_i))}, \quad (2)$$

where the hyperparameters  $\tau$  is a discounting parameter and  $\alpha$  scales the influence of current prediction. This multiplicative update method has a close relation to maximum loss minimization [42] and AdaBoost [15], which can result in improved generalization performance, especially when only a few “rare” samples exist. Moreover, when the gradient update is weighted by the inverse sampling probability, we obtain an unbiased gradient estimation that improves the convergence by reducing its variance [55, 16].

## 4 Adaptive Task Sampling for Meta-Learning

In this section, we first propose the class-based adaptive task sampling method which is a straightforward extension of the instance-based sampling. Then, we discuss its defect and present the class-pair based sampling method. Finally, we propose the greedy class-pair based sampling method, which significantly reduces the computation cost while still generating the identical task distribution as that in the non-greedy approach.

### 4.1 Class-based Sampling

A major challenge of adaptive task sampling for meta-learning is the implicit definition of the task, which is randomly generated by sampling  $K$  classes in each episode. Although direct task based sampling is infeasible, we can adaptively sample classes for each  $K$ -way classification task. With this goal in mind, we propose a **class-based sampling (c-sampling)** approach that updates the class selection probability  $p_C^{t+1}(c)$  in each episode. Given  $S^t$  and  $Q^t$  at episode  $t$ , we could update the class selection probability for each class in current episode  $c \in \mathbb{L}_K^t$  in the following way,

$$p_C^{t+1}(c) \propto (p^t(c))^\tau e^{\alpha \frac{\sum_{(q_n, y_n) \in Q^t} \mathbb{I}[c \neq y_n] p(c|q_n, S^t) + \mathbb{I}[c = y_n] (1 - p(c|q_n, S^t))}{NK}}}. \quad (3)$$

Note that we average the prediction probability of classifying each query sample  $n$  into incorrect classes in  $\mathbb{L}_K^t$ . Then we can sample  $K$  classes without replacement to construct the category set  $\mathbb{L}_K^{t+1}$  for the next episode.

Despite its simplicity, such a sampling approach does suffer from an important limitation. It implicitly assumes that the difficulty of each class is independent. Therefore, it updates the class selection probability in a decoupled way. In concrete words, suppose we have two different tasks: discerning “corgi”, “Akita” and “poodle” and discerning “corgi”, “car” and “people”. Obviously, it is quite hard to tell “corgi” in the first task while it could be easy in the second one. This would be a challenging aspect for updating the class selection probability as the class-based sampling is agnostic to the context of the task and could accidentally assign contradictory scores to the same class. Secondly, even if the class selection probability is updated correctly, it cannot ensure that difficult tasks are generated properly. That is, assembling the most difficult classes do not necessarily lead to a difficult task.

## 4.2 Class-Pair Based Sampling

To address the above issue, we further propose a **class-pair based sampling (cp-sampling)** approach that exploits the pairwise relationships between classes. This idea is commonly used in the multi-class classification that constructs binary classifiers to discriminate between each pair of classes [3], as two-class problems are much easier to solve. Recently, it has also been considered to extract the pairwise relationships between classes for task-dependent fast adaptation in few-shot learning [40]. In this work, we formulate the task selection probability by leveraging the Markov random field [10] over class pairs. Formally, the probability of choosing a category set  $\mathbb{L}_K^{t+1}$  at episode  $t + 1$  is defined as:

$$p_{CP}^{t+1}(\mathbb{L}_K^{t+1}) \propto \prod_{(i,j) \in \mathbb{L}_K^{t+1}} C^t(i,j) \quad \text{s.t. } i, j \in \mathbb{C}_{tr}, \quad (4)$$

where  $C^t(i, j)$  is a potential function over class pair  $(i, j)$  at episode  $t$ . Notice that the classes in  $\mathbb{C}_{tr}$  form a complete and undirected graph. The category set  $\mathbb{L}_K^{t+1}$  that have a relatively high probability to be selected are those  $K$ -cliques with large potentials. Similarly, we adaptively update the potential function  $C^{t+1}(i, j)$  according to

$$C^{t+1}(i, j) \leftarrow (C^t(i, j))^{\tau} e^{\alpha \bar{p}((i,j)|\mathbb{S}^t, \mathbb{Q}^t)}, \quad i \neq j \quad (5)$$

where  $\bar{p}((i, j)|\mathbb{S}^t, \mathbb{Q}^t)$  denotes the average prediction probability that classifies query samples in class  $j$  into its incorrect class  $i$  or vice versa. Specifically, we define it as

$$\bar{p}((i, j)|\mathbb{S}^t, \mathbb{Q}^t) = \frac{\sum_{(q_n, y_n=j) \in \mathbb{Q}^t} p(c=i|q_n, \mathbb{S}^t)}{N} + \frac{\sum_{(q_n, y_n=i) \in \mathbb{Q}^t} p(c=j|q_n, \mathbb{S}^t)}{N}. \quad (6)$$

## 4.3 Greedy Class-Pair Based Sampling

It is important to note that class-pair based sampling has the disadvantage that  $\binom{K}{2} \cdot \binom{|\mathbb{C}_{tr}|}{K}$  multiplication operations need to be performed for calculating  $p_{CP}^{t+1}(\mathbb{L}_K^{t+1})$  for






Fig. 2: A toy example to illustrate how greedy class-pair based sampling chooses 4-class category set  $\mathbb{L}_4^{t+1}$  from 5 classes. The left correlation matrix indicates the class-pair potentials  $C^t$  and the right part denotes the state of each step in sequential sampling. The blue number on the right denotes the chosen class and the red circle highlights the highest unnormalized class selection probability.  $\odot$  denotes the element-wise multiplication.

different combinations of  $K$ -class in the category set. To significantly reduce the complexity, we now design a **greedy class-pair based sampling (gcp-sampling)** method, which samples not only at the cost  $O(K)$  but also from a distribution identical to that in Eq. (4), due to the independence of the potential function  $C^t(i, j)$  over class pairs. In particular, we sequentially sample classes in  $K - 1$  steps based on the previous results. At episode  $t$ , we first sample two classes based on class-pair potential function  $C^t(i, j)$ . Then we iteratively sample a new class based on the already sampled classes. Figure 2 gives an example to illustrate the process. Formally, the task selection probability is defined as

$$p_{GCP}^{t+1}(\mathbb{L}_{k+1}^{t+1}) \propto \begin{cases} C^t(i, j), & k = 1 \\ p(c|\mathbb{L}_k^{t+1}, C^t), & k > 1 \end{cases} \quad (7)$$

where  $p(c = i|\mathbb{L}_k^{t+1}, C^t) \propto \prod_{j \in \mathbb{L}_k^{t+1}} C^t(i, j)$ . It considers the joint probability over class pairs between the chosen class  $i$  and every sampled class  $j$  in the category set  $\mathbb{L}_k^{t+1}$ . Compared to the distribution in Eq. (4), the greedy sampling approach in Eq. (7) has a different normalization constant in each step  $k$ . However, for the evaluation of task selection distribution, the unnormalized joint probability over the class pairs of a specific category set is identical which makes the distribution in Eq. (7) exactly the same as that in Eq. (4), which we prove in Proposition 1.

**Proposition 1** *The greedy class-pair based sampling strategy in Eq. (7) is identical to the class-pair based sampling in Eq. (4).*

*Proof.* We present a proof by induction. It is obvious that  $p_{GCP}^{t+1}(\mathbb{L}_2^{t+1}) = p_{CP}^{t+1}(\mathbb{L}_2^{t+1})$  since  $p_{GCP}^{t+1}(\mathbb{L}_2^{t+1}) \propto C^t(i, j)$ . Now let us consider a general case where we have previously sampled  $k$  classes with  $\mathbb{L}_k^{t+1}$  and are about to sample the  $(k + 1)$ -th class.

**Algorithm 1** gcp-sampling: Greedy Class-Pair based Sampling in K-Way-M-Shot**Require:** meta-training data  $\mathbb{D}_{tr}$ , hyperparameters  $\alpha, \tau, T$ 

- 1: Randomly initialize meta model parameter  $\theta$ . Initialize class-pair potentials  $C$  by ones
- 2: **for**  $t = 1, \dots, T$  **do**
- 3:   Initialize  $\mathbb{L}_0^t$  by an empty set. Initialize  $p(c|\mathbb{L}_0^t, C^{t-1})$  by  $\frac{1}{|\mathbb{C}_{tr}|}$
- 4:   Sample class pair  $(i, j) \propto C(i, j)$ , add class  $i$  and  $j$  into  $\mathbb{L}_0^t$
- 5:   **for**  $k = 2, \dots, K - 1$  **do**
- 6:     Update  $p(c = i|\mathbb{L}_k^t, C^{t-1}) \propto \prod_{j \in \mathbb{L}_k^t} C^{t-1}(i, j)$
- 7:     Sample class  $c$  based on  $p(c|\mathbb{L}_k^t, C^{t-1})$ , add class  $c$  into  $\mathbb{L}_{k+1}^t$
- 8:   **end for**
- 9:   Construct support set  $\mathbb{S}^t$  and query set  $\mathbb{Q}^t$  by sampling  $M$  and  $N$  image per class in category set  $\mathbb{L}_K^t$ , respectively
- 10:   Update meta model  $\theta$  based on support set and query set
- 11:   Update class-pair potentials  $C$  according to Eq. (5)
- 12: **end for**
- 13: **return**  $\theta_T$

Suppose we sample a new class  $l$  to generate  $\mathbb{L}_{k+1}^{t+1}$ , according to Eq. (7), we have

$$\begin{aligned}
 p_{GCP}^{t+1}(\mathbb{L}_{k+1}^{t+1}) &= p_{GCP}^{t+1}(\mathbb{L}_k^{t+1})p(c = l|\mathbb{L}_k^{t+1}, C^t) \propto \prod_{(i,j) \subset \mathbb{L}_k^{t+1}} C^t(i, j) \prod_{j \in \mathbb{L}_k^{t+1}} C^t(l, j) \\
 &= \prod_{(i,j) \subset \mathbb{L}_{k+1}^{t+1}} C^t(i, j) = p_{CP}^{t+1}(\mathbb{L}_{k+1}^{t+1}).
 \end{aligned} \tag{8}$$

The pseudocode of the proposed gcp-sampling algorithm is given in Algorithm 1. Due to the space limitation, we leave the theoretical analysis of the proposed gcp-sampling method in terms of its generalization ability to the supplementary material.

## 5 Experiments

In this section, we evaluate the proposed adaptive task sampling method on two few-shot classification benchmarks: miniImageNet [50] and CIFAR-FS [6]. We first introduce the datasets and settings, and then present a comparison to state-of-the-art methods, followed by a detailed evaluation of the compatibility when integrating with different meta-learning algorithms and the efficacy of different sampling strategies. Finally, we demonstrate qualitative results to characterize the gcp-sampling.

### 5.1 Datasets and Evaluation

**Datasets.** We conduct experiments to evaluate our method on two few-shot classification benchmarks. Firstly, **miniImageNet** [50] is widely used for few-shot learning, which is constructed based on the ImageNet dataset [39] and thus has high diversity and complexity. This dataset has 100 classes with 600  $84 \times 84$  images per class. These classes are divided into 64, 16 and 20 classes for meta-training, meta-validation and

meta-test, respectively, as suggested earlier [38, 12, 46]. Secondly, **CIFAR-FS** is another recent few-shot image classification benchmark [6] constructed by randomly sampling from the CIFAR-100 dataset [23] using the same criteria as the miniImageNet, and has the same number of classes and samples. The limited resolution of  $32 \times 32$  makes the task still difficult. We also use the 64 / 16 / 20 divisions for consistency with previous studies [6, 26].

**Evaluation metric.** We report the mean accuracy (%) of 1000 randomly generated episodes as well as the 95% confidence intervals on the meta-test set. In every episode during meta-test, each class has 15 queries.

## 5.2 Implementation Details

We validate the efficacy of the proposed adaptive sampling strategy on different meta-learning methods, including the gradient-based meta-learning methods: MAML [12], Reptile [36] and MAML++ [4], and metric-based meta-learning methods: PN [44] and MN [50]. We evaluate our adaptive task sampling strategy on all these meta-learning algorithms based on their open-source implementations<sup>1</sup>.

**Network Architectures.** We conduct experiments with 2 different feature extractor architectures, Conv-4 and ResNet-12. **Conv-4** is a shallow embedding function proposed by [50] and widely used [12, 4, 44, 36]. It is composed of 4 convolutional blocks, each of which comprises a 64-filter  $3 \times 3$  convolution, batch normalization (BN) [19], a ReLU nonlinearity and a  $2 \times 2$  max-pooling layer. We also adopt a deep backbone **ResNet-12** [17], which achieves significant improvement in recent works [33, 34, 37]. It consists of 4 residual blocks, each of which has three  $3 \times 3$  convolutional layers and a  $2 \times 2$  max-pooling layer. The number of filters starts from 64 and is doubled every next block. There is also a mean-pooling layer compressing the feature maps to a feature embedding in the end.

In our experiments, we integrate gcp-sampling with PN, MetaOptNet-RR and MetaOptNet-SVM with ResNet-12 to compare with state of the arts. We follow the settings of [26] and use SGD with Nesterov momentum of 0.9 and weight decay of 0.0005. Besides, we use Conv-4 to evaluate the compatibility when integrating with different meta-learning algorithms and the efficacy of different sampling strategies. We follow the settings of [8] and use Adam [22] optimizer with an initial learning rate of 0.001.

## 5.3 Results and Analysis

**Comparison with state-of-the-art.** Tables 1 and 2 present the 5-way 1-shot and 5-way 5-shot results on miniImageNet and CIFAR-FS datasets, respectively. Note that it

<sup>1</sup> Available at the following sites. (a) Matching Network [50]: <https://github.com/wyharveychen/CloserLookFewShot/>, (b) PN [44]: <https://github.com/kjunelee/MetaOptNet>, <https://github.com/wyharveychen/CloserLookFewShot/>, (c) MAML [12] and MAML++ [4]: <https://github.com/AntreasAntoniou/HowToTrainYourMAMLPytorch>, (d) Reptile [36]: <https://github.com/dragen1860/Reptile-Pytorch>.

Table 1: Average 5-way, 1-shot and 5-shot classification accuracies (%) on the miniImageNet dataset. \* denotes the results from [26].

Methods	Backbone	5-way-1-shot	5-way-5-shot
Matching Network [50]	CONV-4	43.44 ± 0.77	55.31 ± 0.73
Relation Network [47]	CONV-4	50.44 ± 0.82	65.32 ± 0.70
PN [44]	CONV-4	49.42 ± 0.78	68.20 ± 0.66
MAML [12]	CONV-4	48.70 ± 1.84	63.11 ± 0.92
MAML++ [4]	CONV-4	52.15 ± 0.26	68.32 ± 0.44
MAML++, AS (ours)	CONV-4	52.34 ± 0.81	69.21 ± 0.68
Bilevel Programming [13]	ResNet-12	50.54 ± 0.85	64.53 ± 0.68
MetaGAN [54]	ResNet-12	52.71 ± 0.64	68.63 ± 0.67
SNAIL [33]	ResNet-12	55.71 ± 0.99	68.88 ± 0.92
adaResNet [34]	ResNet-12	56.88 ± 0.62	71.94 ± 0.57
TADAM [37]	ResNet-12	58.50 ± 0.30	76.70 ± 0.30
MTL [46]	ResNet-12	61.2 ± 1.8	75.5 ± 0.8
PN* [26]	ResNet-12	59.25 ± 0.64	75.60 ± 0.48
PN with gcp-sampling	ResNet-12	<b>61.09</b> ± 0.66	<b>76.80</b> ± 0.49
MetaOptNet-RR [26]	ResNet-12	61.41 ± 0.61	77.88 ± 0.46
MetaOptNet-RR with gcp-sampling	ResNet-12	<b>63.02</b> ± 0.63	<b>78.91</b> ± 0.46
MetaOptNet-SVM [26]	ResNet-12	62.64 ± 0.61	78.63 ± 0.46
MetaOptNet-SVM with gcp-sampling	ResNet-12	<b>64.01</b> ± 0.61	<b>79.78</b> ± 0.47

Table 2: Average 5-way, 1-shot and 5-shot classification accuracies (%) on the CIFAR-FS dataset. \* denotes the results from [26].

Methods	Backbone	5-way-1-shot	5-way-5-shot
Relation Network [47]	CONV-4	55.0 ± 1.0	69.3 ± 0.8
PN* [44]	CONV-4	55.5 ± 0.7	72.0 ± 0.6
MAML* [12]	CONV-4	58.9 ± 1.9	71.5 ± 1.0
GNN [41]	CONV-4	61.9	75.3
R2D2 [26]	CONV-4	65.3 ± 0.2	79.4 ± 0.1
PN* [26]	ResNet-12	72.2 ± 0.7	84.2 ± 0.5
PN with gcp-sampling	ResNet-12	<b>74.1</b> ± 0.7	<b>84.5</b> ± 0.5
MetaOptNet-RR [26]	ResNet-12	72.6 ± 0.7	84.3 ± 0.5
MetaOptNet-RR with gcp-sampling	ResNet-12	<b>74.2</b> ± 0.7	<b>85.1</b> ± 0.4
MetaOptNet-SVM [26]	ResNet-12	72.0 ± 0.7	84.2 ± 0.5
MetaOptNet-SVM with gcp-sampling	ResNet-12	<b>73.9</b> ± 0.7	<b>85.3</b> ± 0.5

shows the highest accuracies for which the iterations are chosen by validation. For our approach, we integrate gcp-sampling with PN, MON-RR and MON-SVM, which are strong baselines. For all cases, we achieve comparable performance surpassing prior methods by a meaningful margin. For example, PN with gcp-sampling outperforms the PN with ResNet-12 by around 1.84 and 1.2 percentage points in miniImageNet and 1.89 and 1.0 percentage points in CIFAR-FS. It is worth noting that the adaptive task

Table 3: Average 5-way classification accuracies (%) on miniImageNet and CIFAR-FS. All methods use shallow feature backbone (Conv-4). <sup>†</sup> denotes the local replication results. We run PN without oversampling the number of ways.

Model	miniImageNet		CIFAR-FS	
	1-shot	5-shot	1-shot	5-shot
Matching Network <sup>†</sup>	48.26 ± 0.76	62.27 ± 0.71	53.14 ± 0.85	68.16 ± 0.76
Matching Network with gcp-sampling	<b>49.61</b> ± 0.77	<b>63.23</b> ± 0.75	<b>54.72</b> ± 0.87	<b>69.28</b> ± 0.74
PN <sup>†</sup>	44.15 ± 0.76	63.89 ± 0.71	54.87 ± 0.72	71.64 ± 0.58
PN with gcp-sampling	<b>47.13</b> ± 0.81	<b>64.75</b> ± 0.72	<b>56.12</b> ± 0.81	<b>72.77</b> ± 0.64
Reptile <sup>†</sup>	46.12 ± 0.80	63.56 ± 0.70	55.86 ± 1.00	71.08 ± 0.74
Reptile with gcp-sampling	<b>47.60</b> ± 0.80	<b>64.56</b> ± 0.69	<b>57.25</b> ± 0.99	<b>71.69</b> ± 0.71
MAML <sup>†</sup>	48.25 ± 0.62	64.09 ± 0.70	56.93 ± 0.99	72.10 ± 0.74
MAML with gcp-sampling	<b>49.65</b> ± 0.85	<b>65.37</b> ± 0.70	<b>57.62</b> ± 0.97	<b>72.51</b> ± 0.72
MAML++ <sup>†</sup>	50.60 ± 0.82	68.24 ± 0.68	58.87 ± 0.97	73.86 ± 0.76
MAML++ with gcp-sampling	<b>52.34</b> ± 0.81	<b>69.21</b> ± 0.68	<b>60.14</b> ± 0.97	<b>73.98</b> ± 0.74

Table 4: Average 5-way classification accuracies (%) on miniImageNet and CIFAR-FS. Using MAML++ on a Conv-4 backbone, we compare different sampling methods: random, c-sampling with hard class, gcp-sampling with hard/uncertain/easy class.

Sampling Strategy	miniImageNet		CIFAR-FS	
	5-way-1-shot	5-way-5-shot	5-way-1-shot	5-way-5-shot
random sampling	50.60 ± 0.82	68.24 ± 0.68	58.87 ± 0.97	73.36 ± 0.76
c-sampling with hard class	51.43 ± 0.75	68.74 ± 0.67	58.61 ± 0.92	73.98 ± 0.72
gcp-sampling with easy class	50.88 ± 0.88	68.22 ± 0.72	58.73 ± 1.14	73.41 ± 0.76
gcp-sampling with uncertain class	51.73 ± 0.87	69.01 ± 0.72	59.43 ± 1.02	73.84 ± 0.82
gcp-sampling with hard class	<b>52.34</b> ± 0.81	<b>69.21</b> ± 0.68	<b>60.14</b> ± 0.97	<b>74.58</b> ± 0.74

sampling method is orthogonal to the meta-learning algorithm. Moreover, even for a deep feature backbone, our approach is still able to preserve the performance gain.

**Compatibility with different meta-learning algorithms** Next, we study the impact of gcp-sampling when integrating with different types of meta-learning algorithm. We consider gradient-based meta-learning methods: MAML, Reptile and MAML++, and metric-based meta-learning methods: PN and MN. The results in Table 3 demonstrate that using gcp-sampling for meta-learning methods consistently improves the few-shot classification performance. Moreover, the performance improvement is more significant for 1-shot classification than 5-shot classification.

**Efficacy of different adaptive task sampling strategies.** In literature, there exist contradicting ideas in adaptive sampling strategies which work well in different scenarios [7]. Preferring easier samples may be effective when solving challenging problems

containing noise or outliers. The opposite hard sample mining strategy may improve the performance since it is more likely to be minority classes. Therefore, we explore different sampling strategies for meta-learning for few-shot classification. As defined in Eq. (6) for hard class, the probability of easy class is  $1 - \bar{p}(i, j)$  and uncertain class is  $(1 - \bar{p}(i, j))(\bar{p}(i, j))$ , respectively. We report the results in Table 4. We observe that gcp-sampling with hard or uncertain class outperforms that with random sampling, but uncertain sampling offers a smaller improvement. We also compare gcp-sampling with c-sampling, in which c-sampling achieves similar performance as random sampling, verifying the efficacy of using class pairs to represent task difficulty.

**Impact of Hyperparameters  $\alpha$  and  $\tau$**  In the proposed gcp-sampling, the hyperparameter  $\alpha$  controls the aggressiveness of the update while the hyperparameter  $\tau$  controls the degree of forgetting past updates. Here we adopt PN with **ResNet-12** backbone and report the effect of  $\alpha$  and  $\tau$  on the testing performance in Figure 3.




Fig. 3: Impact of hyperparameters  $\alpha$  and  $\tau$ . First row (a-d): we fix the discounting factor  $\tau = 0.5$  and tune the updating factor  $\alpha$ ; Second row (e-g): we fix  $\alpha = 1$  and tune  $\tau$ .

**Time Cost Analysis** Table 5 shows the time cost comparison between random sampling and gcp-sampling. We adopt PN on the CIFAR-FS dataset and report the average training time for each epoch, which includes task sampling, forward and backward propagation phases. We find that the time taken by gcp-sampling is comparable to the time taken by random-sampling. This is because the training time is dominated by the forward pass and backward pass and the cost of task generation and class-pair potential update is relatively small. Besides, using a deeper backbone significantly increases the time cost but reduces the ratio between gcp-sampling and random-sampling, since it only affects the forward pass and backward pass. Finally, increasing the number of ways would increase the time cost while increasing the number of shots will not. This is because the complexity of gcp-sampling scales linearly to the number of ways.

**Visual analysis of adaptive task sampling.** To qualitatively characterize adaptive task sampling, we visualize the prototype of each class generated by the training procedure

Table 5: Time cost comparison between random sampling and gcp-sampling. All the experiments are conducted with PN on the CIFAR-FS dataset.

	random sampling	gcp-sampling	factor
5-way-1-shot, Conv-4	235.4	251.8	1.070
5-way-1-shot, ResNet-12	531.2	554.6	1.044
5-way-5-shot, Conv-4	342.2	367.3	1.073
5-way-10-shot, Conv-4	471.4	491.0	1.042
5-way-15-shot, Conv-4	617.2	634.6	1.028
10-way-1-shot, Conv-4	411.3	451.7	1.098
15-way-1-shot, Conv-4	624.9	723.5	1.158
20-way-1-shot, Conv-4	816.8	992.5	1.215

of PN with gcp-sampling and random sampling. We use the t-SNE [32] method to convert the prototypes into two-dimensional vectors by preserving the cosine similarity between them. As shown in Figure 4, the classes sampled by random sampling achieve better clustering results than gcp-sampling. This is because gcp-sampling tends to sample classes with highly overlapping embeddings, which is much more difficult to learn for meta-learner.




Fig. 4: Feature embedding of the classes sampled by (a) random sampling and (b) task adaptive sampling. The dimension reduction is performed based on all 64 training classes of CIFAR-FS, while we show only the 5 selected classes in each sub-figure for better visualization.

We also visualize the class-pair potentials constructed by gcp-sampling in Figure 5. We show 16 classes of CIFAR-FS, where the green and red colors denote the classes sampled by random sampling and gcp-sampling, respectively. We can see that the classes sampled by random sampling are often easier to distinguish, which leads to inefficient training, while the gcp-sampling tends to sample the classes that, when combined with other classes, display greater difficulty. We also randomly select some sampled images from each class for observation. As shown in Figure 6, the classes sampled by random sampling do vary greatly (e.g., with unique shapes or colors) and are easier to recognize, while the classes sampled by gcp-sampling are visually more confusing (e.g., small animals or insects in the wild) and much more difficult to distinguish.

Fig. 5: Correlation matrix w.r.t. class-pair potentials. Each element indicates the class-pair potential. The higher the correlation weight (i.e., the darker the color), the higher the probability of this two-class combination being sampled. The green and red colors denote the classes sampled by random sampling and adaptive sampling, respectively.

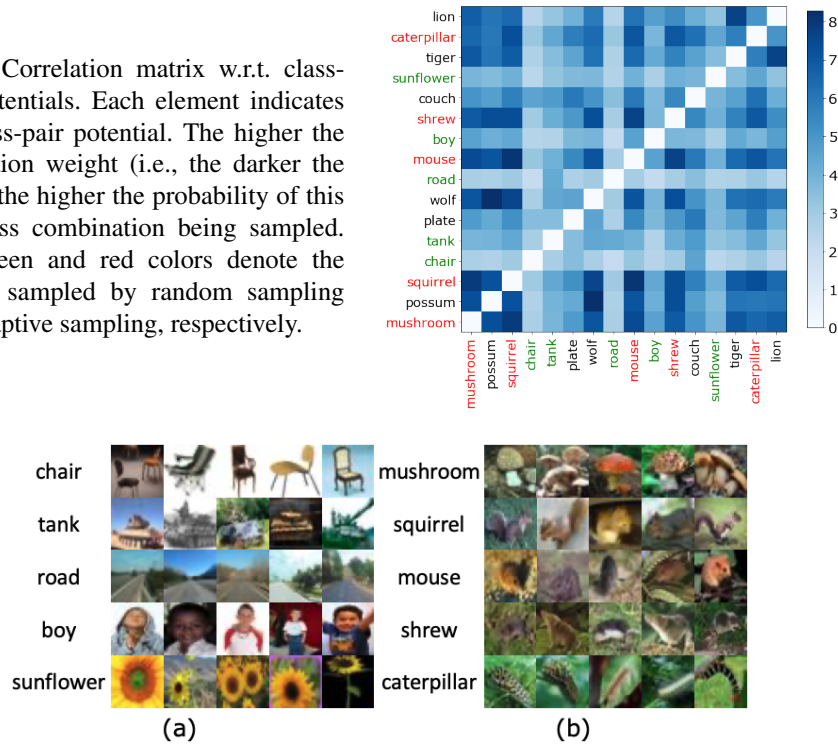


Fig. 6: Sample images from classes by (a) random sampling and (b) gcp-sampling.

## 6 Conclusion

In this paper, we presented an adaptive task sampling method for meta-learning. Our results demonstrated that in meta-learning it is essential for the sampling process to be dependent on tasks, and the proposed method naturally models and exploits this dependence. We showed that the greedy class-pair based sampling method, integrated with PN, MetaOptNet-RR or MetaOptNet-SVM, could achieve competitive results. Furthermore, we demonstrated consistent improvement when integrating the proposed sampling method with different meta-learning methods. Finally, we explore and evaluate different sampling strategies for gcp-sampling, in which the hard class strategy consistently leads to more accurate results.

### Acknowledgment

This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG-RP-2018-001). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.



## References

1. Alain, G., Lamb, A., Sankar, C., Courville, A., Bengio, Y.: Variance reduction in sgd by distributed importance sampling. arXiv preprint arXiv:1511.06481 (2015)
2. Allen-Zhu, Z., Qu, Z., Richtárik, P., Yuan, Y.: Even faster accelerated coordinate descent using non-uniform sampling. In: International Conference on Machine Learning. pp. 1110–1119 (2016)
3. Aly, M.: Survey on multiclass classification methods. *Neural Netw* **19**, 1–9 (2005)
4. Antoniou, A., Edwards, H., Storkey, A.: How to train your maml. arXiv preprint arXiv:1810.09502 (2018)
5. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th annual international conference on machine learning. pp. 41–48. ACM (2009)
6. Bertinetto, L., Henriques, J.F., Torr, P.H., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. arXiv preprint arXiv:1805.08136 (2018)
7. Chang, H.S., Learned-Miller, E., McCallum, A.: Active bias: Training more accurate neural networks by emphasizing high variance samples. In: Advances in Neural Information Processing Systems. pp. 1002–1012 (2017)
8. Chen, W., Liu, Y., Kira, Z., Wang, Y.F., Huang, J.: A closer look at few-shot classification. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019 (2019), <https://openreview.net/forum?id=HkxLXnAcFQ>
9. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014)
10. Cross, G.R., Jain, A.K.: Markov random field texture models. *IEEE Transactions on Pattern Analysis & Machine Intelligence* **PAMI-5**(1), 25–39 (1983)
11. Csiba, D., Richtárik, P.: Importance sampling for minibatches. *The Journal of Machine Learning Research* **19**(1), 962–982 (2018)
12. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: Proceedings of the 34th International Conference on Machine Learning-Volume 70. pp. 1126–1135. JMLR. org (2017)
13. Franceschi, L., Frasconi, P., Salzo, S., Grazzi, R., Pontil, M.: Bilevel programming for hyperparameter optimization and meta-learning. In: International Conference on Machine Learning. pp. 1563–1572 (2018)
14. Freund, Y., Schapire, R.: A short introduction to boosting. *Journal-Japanese Society For Artificial Intelligence* **14**(771-780), 1612 (1999)
15. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* **55**(1), 119–139 (1997)
16. Gopal, S.: Adaptive sampling for sgd by exploiting side information. In: International Conference on Machine Learning. pp. 364–372 (2016)
17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
18. Horváth, S., Richtárik, P.: Nonconvex variance reduced optimization with arbitrary sampling. arXiv preprint arXiv:1809.04146 (2018)
19. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International Conference on Machine Learning. pp. 448–456 (2015)
20. Katharopoulos, A., Fleuret, F.: Biased importance sampling for deep neural network training. arXiv preprint arXiv:1706.00043 (2017)

21. Katharopoulos, A., Fleuret, F.: Not all samples are created equal: Deep learning with importance sampling. arXiv preprint arXiv:1803.00942 (2018)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
23. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
24. Lake, B.M., Salakhutdinov, R., Tenenbaum, J.B.: Human-level concept learning through probabilistic program induction. *Science* **350**(6266), 1332–1338 (2015)
25. Landau, B., Smith, L.B., Jones, S.S.: The importance of shape in early lexical learning. *Cognitive development* **3**(3), 299–321 (1988)
26. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10657–10665 (2019)
27. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few-shot learning. arXiv preprint arXiv:1707.09835 (2017)
28. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
29. Liu, L., Zhou, T., Long, G., Jiang, J., Zhang, C.: Learning to propagate for graph meta-learning. arXiv preprint arXiv:1909.05024 (2019)
30. London, B.: A pac-bayesian analysis of randomized learning with application to stochastic gradient descent. In: Advances in Neural Information Processing Systems. pp. 2931–2940 (2017)
31. Loshchilov, I., Hutter, F.: Online batch selection for faster training of neural networks. arXiv preprint arXiv:1511.06343 (2015)
32. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
33. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. In: ICLR (2017)
34. Munkhdalai, T., Yuan, X., Mehri, S., Trischler, A.: Rapid adaptation with conditionally shifted neurons. In: International Conference on Machine Learning. pp. 3661–3670 (2018)
35. Naik, D.K., Mammone, R.J.: Meta-neural networks that learn by learning. In: [Proceedings 1992] IJCNN International Joint Conference on Neural Networks. vol. 1, pp. 437–442. IEEE (1992)
36. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999 (2018)
37. Oreshkin, B., López, P.R., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. In: Advances in Neural Information Processing Systems. pp. 721–731 (2018)
38. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2016)
39. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International journal of computer vision* **115**(3), 211–252 (2015)
40. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019 (2019), <https://openreview.net/forum?id=BJgklhAcK7>
41. Satorras, V.G., Bruna, J.: Few-shot learning with graph neural networks. In: ICLR (2018)
42. Shalev-Shwartz, S., Wexler, Y.: Minimizing the maximal loss: How and why. In: ICML. pp. 793–801 (2016)

43. Shrivastava, A., Gupta, A., Girshick, R.: Training region-based object detectors with online hard example mining. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 761–769 (2016)
44. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems. pp. 4077–4087 (2017)
45. Song, H., Kim, S., Kim, M., Lee, J.G.: Ada-boundary: Accelerating the dnn training via adaptive boundary batch selection (2018)
46. Sun, Q., Liu, Y., Chua, T.S., Schiele, B.: Meta-transfer learning for few-shot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 403–412 (2019)
47. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1199–1208 (2018)
48. Thrun, S., Pratt, L.: Learning to learn: Introduction and overview. In: Learning to learn, pp. 3–17. Springer (1998)
49. Triantafillou, E., Zhu, T., Dumoulin, V., Lamblin, P., Xu, K., Goroshin, R., Gelada, C., Swersky, K., Manzagol, P.A., Larochelle, H.: Meta-dataset: A dataset of datasets for learning to learn from few examples. arXiv preprint arXiv:1903.03096 (2019)
50. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: Advances in neural information processing systems. pp. 3630–3638 (2016)
51. Ze, H., Senior, A., Schuster, M.: Statistical parametric speech synthesis using deep neural networks. In: 2013 IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 7962–7966. IEEE (2013)
52. Zhang, C., Kjellstrom, H., Mandt, S.: Determinantal point processes for mini-batch diversification. arXiv preprint arXiv:1705.00607 (2017)
53. Zhang, C., Öztireli, C., Mandt, S., Salvi, G.: Active mini-batch sampling using repulsive point processes. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 5741–5748 (2019)
54. Zhang, R., Che, T., Ghahramani, Z., Bengio, Y., Song, Y.: Metagan: An adversarial approach to few-shot learning. In: Advances in Neural Information Processing Systems. pp. 2365–2374 (2018)
55. Zhao, P., Zhang, T.: Stochastic optimization with importance sampling for regularized loss minimization. In: international conference on machine learning. pp. 1–9 (2015)