

Meta-Learning with Network Pruning

Hongduan Tian¹, Bo Liu², Xiao-Tong Yuan¹, and Qingshan Liu¹

¹ B-DAT Lab, Nanjing University of Information Science and Technology, Nanjing, 210044, China

² JD Finance America Corporation, Mountain View, CA 94043, USA
{hongduan.tian,kfliubo,xtyuan1980}@gmail.com, qsliu@nuist.edu.cn

Abstract. *Meta-learning* is a powerful paradigm for few-shot learning. Although with remarkable success witnessed in many applications, the existing optimization based meta-learning models with over-parameterized neural networks have been evidenced to overfit on training tasks. To remedy this deficiency, we propose a network pruning based meta-learning approach for overfitting reduction via explicitly controlling the capacity of network. A uniform concentration analysis reveals the benefit of network capacity constraint for reducing generalization gap of the proposed meta-learner. We have implemented our approach on top of Reptile assembled with two network pruning routines: Dense-Sparse-Dense (DSD) and Iterative Hard Thresholding (IHT). Extensive experimental results on benchmark datasets with different over-parameterized deep networks demonstrate that our method not only effectively alleviates meta-overfitting but also in many cases improves the overall generalization performance when applied to few-shot classification tasks.

Keywords: Meta-Learning; Few-shot Learning; Network Pruning; Sparsity; Generalization Analysis.

1 Introduction

The ability of adapting to a new task with several trials is essential for artificial agents. The goal of few-shot learning [25] is to build a model which is able to get the knack of a new task with limited training samples. Meta-learning [26, 3, 30] provides a principled way to cast few-shot learning as the problem of *learning-to-learn*, which typically trains a hypothesis or learning algorithm to memorize the experience from previous tasks for a future task learning with very few samples. The practical importance of meta-learning has been witnessed in many vision and online/reinforcement learning applications including image classification [23, 17], multi-arm bandit [29] and 2D navigation [5].

Among others, one particularly simple yet successful meta-learning paradigm is first-order optimization based meta-learning which aims to train hypotheses that can quickly adapt to unseen tasks by performing one or a few steps of (stochastic) gradient descent [23, 5]. Reasons for the recent increasing attention to this class of gradient-optimization based methods include their outstanding efficiency and scalability exhibited in practice [21].

Challenge and motivation. A challenge in the existing meta-learning approaches is their tendency to overfit [20, 33]. When training an over-parameterized meta-learner such as very deep and/or wide convolutional neural networks (CNN) which are powerful for representation learning, there are two sources of potential overfitting at play: the inter-task overfitting of meta-learner (or *meta-overfitting*) to the training tasks and the inner-task overfitting of task-specific learner to the task training data. There have been recent efforts put to deal with inner-task overfitting [16, 36]. The study on the inter-task meta-overfitting, however, still remains under explored. Since in principle the optimization-based meta-learning is designed to learn fast from small amount of data in new tasks, we expect the meta-overfitting to play a more important role in influencing the overall generalization performance of the trained meta-learner.

Sparsity model is a promising tool for high-dimensional machine learning with guaranteed statistical efficiency and robustness to overfitting [19, 34, 1]. It has been theoretically and numerically justified by [2] that sparsity benefits considerably the generalization performance of deep neural networks. In the regime of compact deep learning, the so called *network pruning* technique has been widely studied and evidenced to work favorably in generating sparse subnetworks without compromising generalization performance [6, 11, 7]. Inspired by these remarkable success of sparsity models, it is natural to conjecture that sparsity would also be beneficial for enhancing the robustness of optimization based meta-learning to meta-overfitting.

Our contribution. In this paper, we present a novel gradient-based meta-learning approach with explicit network capacity constraint for overfitting reduction. The problem is formulated as learning a sparse meta-initialization network from training tasks such that in a new task the learned subnetwork can quickly converge to the optimal solution via gradient descent. The core idea is to reduce meta-overfitting by controlling the counts of the non-zero parameters in the meta-learner during the training phase. Theoretically, we have established a uniform generalization gap bound for the proposed sparse meta-learner showing the benefit of capacity constraint for improving its generalization performance. Practically, we have implemented our approach in a joint algorithmic framework of Reptile [21] with network pruning, along with two instantiations using Dense-Sparse-Dense (DSD) [7] and Iterative Hard Thresholding (IHT) [11] as network pruning subroutines, respectively. The actual performance of our approach has been extensively evaluated on few-shot classification tasks with over-parameterized wide CNNs. The obtained results demonstrate that our method can effectively alleviate overfitting and achieve similar or even superior generalization performance to the conventional dense models.

2 Related Work

Optimization-based meta-learning. The family of optimization-based meta-learning approaches usually learn a good hypothesis which can be fast adapted to unseen tasks [23, 5, 21, 12]. Compared to the metric [13, 27] and memory [32, 25]

based meta-learning algorithms, optimization based meta-learning algorithms are gaining increasing attention due to their simplicity, versatility and effectiveness. As a recent leading framework for optimization-based meta-learning, MAML [5] is designed to estimate a meta-initialization network which can be well fine-tuned in an unseen task via only one or few steps of minibatch gradient descent. Although simple in principle, MAML requires computing Hessian-vector product for back-propagation, which could be computationally expensive when the model is big. The first-order MAML (FOMAML) is therefore proposed to improve the computational efficiency by simply ignoring the second-order derivatives in MAML. Reptile [21] is another approximated first-order algorithm which works favorably since it maximizes the inner product between gradients from the same task yet different minibatches, leading to improved model generalization. Recently, several hypothesis biased regularized meta-learning approaches have been studied in [4, 12, 35] with provable strong generalization performance guarantees provided for convex problems. In [16], the meta-learner is treated as a feature embedding module of which the output is used as input to train a multi-class kernel support vector machine as base learner. To deal with overfitting, the CAVIA method [36] decomposes the meta-parameters into the so called context parameters and shared parameters. The context parameters are updated for task adaption with limited capacity while the shared parameters are meta-trained for generalization across tasks.

Network pruning. Early network weight pruning algorithms date back to Optimal Brain Damage [15] and Optimal Brain Surgeon [10]. A dense-to-sparse algorithm was developed by [8] to firstly remove near-zero weights and then fine tune the preserved weights. As a serial work of dense-to-sparse, the dense-sparse-dense (DSD) method [7] was proposed to re-initialize the pruned parameters as zero and retrain the entire network after the dense-to-sparse pruning phase. The iterative hard thresholding (IHT) method [11] shares a similar spirit with DSD to conduct multiple rounds of iteration between pruning and retraining. [28] proposed a data-free method to prune the neurons in a trained network. In [18], an L_0 -norm regularized risk minimization framework was proposed to learn sparse networks during training. More recently, [6] introduced and studied the “lottery ticket hypothesis” which assumes that once a network is initialized, there should exist an optimal subnetwork, which can be learned by pruning, that performs as well as the original network or even superior.

Despite the remarkable success achieved by both meta-learning and network pruning, it still remains largely open to investigate the impact of network pruning on alleviating the meta-overfitting of optimization based meta-learning, which is of primal interest to our study in this paper.

3 Method

3.1 Problem Setup

We consider the N -way K -shot problem as defined in [31]. Tasks are sampled from a specific distribution $p(\mathcal{T})$ and will be divided into *meta training set* \mathcal{S}^{tr} ,

meta validation set \mathcal{S}^{val} , and meta testing set \mathcal{S}^{test} . Classes in different datasets are disjoint (i.e., the class in \mathcal{S}^{tr} will not appear in \mathcal{S}^{test}). During training, each task is made up of support set \mathcal{D}^{supp} and query set \mathcal{D}^{query} . Both \mathcal{D}^{supp} and \mathcal{D}^{query} are sampled from the same classes of \mathcal{S}^{tr} . \mathcal{D}^{supp} is used for training while \mathcal{D}^{query} is used for evaluation. For a N -way K -shot classification task, we sample N out of the C classes from dataset, and then K samples are sampled from each of these classes to form \mathcal{D}^{supp} , namely $\mathcal{D}^{supp} = \{(\mathbf{x}_c^k, y_c^k), k = 1, 2, \dots, K; c = 1, 2, \dots, N\}$. For example, for a 5-way 2-shot task, we sample 2 data-label pairs from each of 5 classes, thus, such a task has 10 samples. Usually, several other samples of the same classes will be sampled to compose \mathcal{D}^{query} . For example, \mathcal{D}^{query} is used in Reptile [21] in evaluation steps. We use the loss function $\ell(v, y)$ to measure the discrepancy between the predicted score vector $v \in \mathbb{R}^C$ and the true label $y \in \{1, \dots, C\}$.

Notation. For an integer n , we denote $[n]$ as the abbreviation of the index set $\{1, \dots, n\}$. We use \odot to denote the element-wise product operator. We say a function $g: \mathbb{R}^p \mapsto \mathbb{R}$ is G -Lipschitz continuous if $|g(\theta) - g(\theta')| \leq G\|\theta - \theta'\|_2$, and g is H -smooth if it obeys $\|\nabla g(\theta) - \nabla g(\theta')\|_2 \leq H\|\theta - \theta'\|_2$.

3.2 Meta-Learning with Model Capacity Constraint

Our ultimate goal is to learn a good initialization of parameters for a convolutional neural network $f_\theta: \mathcal{X} \mapsto \mathcal{Y}$, where θ is the model parameters set, from a set of training tasks such that the learned initialization network generalizes well to future unseen tasks. Inspired by the recent remarkable success of MAML [5] and the strong generalization capability of sparse deep learning models [6, 2], during sparse(or network pruning) phase, we propose to learn from previous task experience a sparse subnetwork started from which the future task-specific networks can be efficiently learned using first-order optimization methods. To this end, we introduce the following layer-wise sparsity constrained stochastic first-order meta-learning formulation:

$$\min_{\theta} \mathcal{R}(\theta) := \mathbb{E}_{T \sim p(\mathcal{T})} \left[\mathcal{L}_{\mathcal{D}_T^{query}} \left(\theta - \eta \nabla_{\theta} \mathcal{L}_{\mathcal{D}_T^{supp}}(\theta) \right) \right], \text{ s.t. } \|\theta_l\|_0 \leq k_l, l \in [L], \quad (1)$$

where $\mathcal{L}_{\mathcal{D}_T^{supp}}(\theta) = \frac{1}{NK} \sum_{(\mathbf{x}_c^k, y_c^k) \in \mathcal{D}_T^{supp}} \ell(f_\theta(\mathbf{x}_c^k), y_c^k)$ is the empirical risk for task T and $\mathcal{L}_{\mathcal{D}_T^{query}}(\theta)$ is similarly defined as the loss evaluated over the query set and η is the learning rate. In the constraint, $\|\theta_l\|_0$ denotes the number of non-zero entries in the parameters of l -th layer θ_l which is required to be no larger than a user-specified sparsity level k_l , and L is the total number of network layers.

In general, the mathematical formulation of task distribution $p(\mathcal{T})$ is unknown but we usually have access to a set of i.i.d. training tasks $S = \{T_i\}_{i=1}^M$ sampled from $p(\mathcal{T})$. Thus the following empirical version of the population form in equation 1 is alternatively considered for training:

$$\min_{\theta} \mathcal{R}_S(\theta) := \frac{1}{M} \sum_{i=1}^M \left[\mathcal{L}_{\mathcal{D}_{T_i}^{query}} \left(\theta - \eta \nabla_{\theta} \mathcal{L}_{\mathcal{D}_{T_i}^{supp}}(\theta) \right) \right], \text{ s.t. } \|\theta_l\|_0 \leq k_l, l \in [L]. \quad (2)$$

To compare with MAML, our model shares an identical objective function, but with the layer-wise sparsity constraints $\|\theta_l\|_0 \leq k_l$ imposed for the purpose of

enhancing learnability of the over-parameterized meta-initialization network. In view of the “lottery ticket hypothesis” [6], the model in equation 2 can be interpreted as a first-order meta-learner for estimating a subnetwork, or a “*winning ticket*”, for future task learning. Inspired by the strong statistical efficiency and generalization guarantees of sparsity models [34, 2], we will very shortly show that such a subnetwork is able to achieve advantageous generalization performance over the dense initialization networks learned by vanilla MAML.

3.3 Generalization Analysis

We provide in this section a task-level generalization performance analysis for the proposed model in equation 2. Let p be the total number of parameters in the over-parameterized network and $\Theta \subseteq \mathbb{R}^p$ be the domain of interest for θ . Let $k = \sum_{l=1}^L k_l$ be the total desired sparsity level of the subnetwork. The following uniform concentration bound is our main result.

Theorem 1. *Assume that the domain of interest Θ is bounded by R and the loss function $\ell(f_\theta(\mathbf{x}), y)$ is G -Lipschitz continuous and H -smooth with respect to θ . Suppose that $0 \leq \ell(f_\theta(\mathbf{x}), y) \leq B$ for all pairs $\{f_\theta(\mathbf{x}), y\}$. Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the random draw of S , the generalization gap is uniformly upper bounded for all θ satisfying $\|\theta_l\|_0 \leq k_l, l \in [L]$ as*

$$|\mathcal{R}(\theta) - \mathcal{R}_S(\theta)| \leq \mathcal{O} \left(B \sqrt{\frac{k \log(p\sqrt{M}GR(1 + \eta H)/(Bk)) + \log(1/\delta)}{M}} \right).$$

In comparison to the $\mathcal{O}(\sqrt{p/M})$ uniform bound established in Lemma 1 (see Appendix A.1) for dense networks, the uniform bound established in Theorem 1 is substantially stronger when $k \ll p$, which shows the benefit of network capacity constraint for generalization.

Specially for margin-based multiclass classification, let us consider the margin operator $\mathcal{M}(v, y) := \max_j [v]_j - [v]_y$ associated with the score prediction vector $v \in \mathbb{R}^C$ and label $y \in \{1, \dots, C\}$. Let $\ell_\gamma(f_\theta(\mathbf{x}), y) = h_\gamma(\mathcal{M}(f_\theta(\mathbf{x}), y))$ be a *surrogate* loss of the binary loss (i.e., $\mathbb{1}[y \neq \arg \max_j [f_\theta(\mathbf{x})]_j]$) defined with respect to proper γ -margin based loss h_γ such as the hinge/ramp losses and their smoothed variants [22]. By definition, we must have $\mathbb{1}[y \neq \arg \max_j [f_\theta(\mathbf{x})]_j] \leq \ell_\gamma(f_\theta(\mathbf{x}), y)$. In this case, we denote $\mathcal{R}_{\gamma, S}$ the meta-training risk with loss function ℓ_γ and $\tilde{\mathcal{R}}_\gamma$ the corresponding population risk in which the task-level query loss $\mathcal{L}_{\mathcal{D}_T^{\text{query}}}$ is evaluated using binary loss as classification error. Then as a direct consequence of Theorem 1, we can establish the following result for margin-based prediction.

Corollary 1. *Suppose that the margin-based loss ℓ_γ is used for model training. Then under the conditions in Theorem 1, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ the following bound holds for all θ satisfying $\|\theta_l\|_0 \leq k_l, l \in [L]$:*

$$\tilde{\mathcal{R}}_\gamma(\theta) \leq \mathcal{R}_{\gamma, S}(\theta) + \mathcal{O} \left(B \sqrt{\frac{k \log(p\sqrt{M}GR(1 + \eta H)/(Bk)) + \log(1/\delta)}{M}} \right).$$

Remark 1. We comment that the above $\mathcal{O}(\sqrt{k/M})$ margin bound derived in the context of sparse meta-learning can be readily extended to sparse deep nets training. Also, the bound can be easily generalized for arbitrary convex surrogates (e.g., cross-entropy loss) of binary loss under proper regularity conditions.

4 Algorithm

We have implemented the proposed model in Equation 2 based on Reptile [21] (see Algorithm 2) which is a scalable method for optimization-based meta-learning in form of Equation 2 but without layer-wise sparsity constraint. In order to handle the sparsity constraint, we follow the principles behind the widely applied *dense-sparse-dense* (DSD) [7] and *iterative hard thresholding* (IHT) [11] network pruning algorithms to alternate the Reptile iteration between pruning insignificant weights in each layer and retraining the pruned network.

4.1 Main Algorithm: Reptile with Iterative Network Pruning

The algorithm of our network-pruning-based Reptile method is outlined in Algorithm 1. The learning procedure contains a pre-training phase followed by an iterative procedure of network pruning and retraining. We would like to stress that since our ultimate goal is not to do network compression, but to reduce meta-overfitting via controlling the sparsity level of the meta-initialization network, the final output of our algorithm is typically dense after the retraining phase, which has been evidenced in practice to be effective for improving the generalization performance during testing phase. In the following subsections, we describe the key components of our algorithm in details.

Model Pretraining For model pre-training, we run a few number of Reptile iteration rounds to generate a relatively good initialization. In each loop of the Reptile iteration, we first sample a mini-batch of meta-tasks $\{T_i\}_{i=1}^s$ from the task distribution $p(\mathcal{T})$. Then for each task T_i , we compute the adapted parameters via (stochastic) gradient descent as $\tilde{\theta}_{T_i} = \theta^{(0)} - \eta \nabla_{\theta} \mathcal{L}_{\mathcal{D}_{T_i}^{supp}}(\theta^{(0)})$, where $\tilde{\theta}_{T_i}$ denotes the task-specific parameters learned from each task T_i , $\theta^{(0)}$ is the current initialization of model parameters, η is the inner-task learning rate, and $\mathcal{D}_{T_i}^{supp}$ denotes the support set of task T_i . When all the task-specific parameters are updated, the initialization parameters will be updated according to $\theta^{(0)} = \theta^{(0)} + \beta \left(\frac{1}{s} \sum_{i=1}^s \tilde{\theta}_{T_i} - \theta^{(0)} \right)$ with learning rate β . Here we follow Reptile to use $\frac{1}{s} \sum_{i=1}^s \tilde{\theta}_{T_i} - \theta^{(0)}$ as an approximation to the negative meta-gradient, which has been evidenced to be effective for scaling up the MAML-type first-order meta-learning models [21].

Iterative Network Pruning and Retraining After model pre-training, we proceed to the main loop of our Algorithm 1 that carries out iterative network pruning and retraining.

Algorithm 1: Reptile with Iterative Network Pruning

```

Input : inner loop learning rate  $\eta$ , outer loop learning rate  $\beta$ , layer-wise
          sparsity level  $\{k_l\}_{l=1}^L$ , mini-batch batch size  $s$  for meta training.
Output:  $\theta^{(t)}$ .
Initialization Randomly initialize  $\theta^{(0)}$ .
/* Pre-training with Reptile */
while the termination condition is not met do
  |  $\theta^{(0)} = \text{Reptile}(\theta^{(0)}, \eta, \beta, s)$ ;
end
for  $t = 1, 2, \dots$  do
  /* Pruning phase */
  Generate a network zero-one mask  $\mathcal{M}^{(t)}$  whose non-zero entries at each
  layer  $l$  are those top  $k_l$  entries in  $\theta_l^{(t)}$ ;
  Compute  $\theta_{\mathcal{M}}^{(t)} = \theta^{(t)} \odot \mathcal{M}^{(t)}$ ;
  /* Subnetwork fine-tune with Reptile */
  while the termination condition is not met do
    |  $\theta^{(t)} = \text{Reptile}(\theta_{\mathcal{M}}^{(t)}, \eta, \beta, s)$ ;
  end
  /* Retraining phase */
  while the termination condition is not met do
    |  $\theta^{(t)} = \text{Reptile}(\theta^{(t)}, \eta, \beta, s)$ ;
  end
end

```

Pruning phase. In this phase, we first greedily truncate out of the model a portion of near-zero parameters which are unlikely to contribute significantly to the model performance. To do so, we generate a network binary mask $\mathcal{M}^{(t)}$ whose non-zero entries at each layer l are those top k_l (in magnitude) entries in $\theta_l^{(t)}$, and compute $\theta_{\mathcal{M}}^{(t)} = \theta^{(t)} \odot \mathcal{M}^{(t)}$ as the sparsity restriction of $\theta^{(t)}$. Then we fine-tune the subnetwork over the mask $\mathcal{M}^{(t)}$ by applying Reptile restrictively to this subnetwork with initialization $\theta_{\mathcal{M}}^{(t)}$. Our numerical experiment suggests that sufficient steps of subnetwork fine-tuning tends to substantially improve the stability and convergence behavior of the method.

The fine-tuned subnetwork $\theta_{\mathcal{M}}^{(t)}$ at the end of the pruning phase is expected to reduce the chance of overfitting to noisy data. However, it is also believed that such subnetwork will reduce the capacity of the network, which could in turn lead to potentially biased learning with higher training loss. To remedy this issue, inspired by the retraining trick introduced in [7] for network pruning, we propose to restore the pruned weights that would be beneficial for enhancing the model representation power to improve the overall generalization performance.

Retraining phase. In this phase, the layer-wise sparsity constraints are removed and the pruned parameters are re-activated for fine-tuning. The re-training procedure is almost identical to the pre-training phase, but with the main difference that the former is initialized with the subnetwork generated by

Algorithm 2: Reptile Algorithm [21]

Input : model parameters ϕ , inner loop learning rate η , outer loop learning rate β , mini-batch batch size s for meta training.

Output: the updated ϕ .

Sample a mini-batch tasks $\{T_i\}_{i=1}^s$ of size s ;

For each task T_i , compute the task-specific adapted parameters using gradient descent:

$$\tilde{\phi}_{T_i} = \phi - \eta \nabla_{\phi} \mathcal{L}_{\mathcal{D}_{T_i}^{supp}}(\phi);$$

Update the parameters: $\phi = \phi + \beta \left(\frac{1}{s} \sum_{i=1}^s \tilde{\phi}_{T_i} - \phi \right)$.

the pruning phase while the latter uses random initialization. Such a retraining operation restores the representation capacity of the pruned parameters, which tends to lead to improved generalization performance in practice. For theoretical justification, roughly speaking, since the sparse meta-initialization network obtained in the pruning phase generalizes well in light of Theorem 1, it is expected to serve as a good initialization for future retraining via gradient descent. Then according to the stability theory of gradient descent methods [9], the output dense network will also generalize well if the retraining phase converges quickly.

4.2 Two Substantialized Implementations

Reptile with DSD pruning. The DSD method is an effective network pruning approach for preventing the learned model from capturing noise during the training [7]. By implementing the main loop with $t = 1$, the proposed Algorithm 1 reduces to a DSD-based Reptile method for first-order meta-learning.

Reptile with IHT pruning. The IHT method [11] is another representative network pruning approach which shares a similar dense-sparse-dense spirit with DSD. Different from the one-shot weight pruning and network training by DSD, IHT is designed to perform multiple rounds of iteration between pruning and retraining, and hence is expected to have better chance to find an optimal sparse subnetwork than DSD does. By implementing the main loop of Algorithm 1 with $t > 1$, we actually obtain a variant of Reptile with IHT-type network pruning.

5 Experiments

In this section, we carry out a numerical study for algorithm performance evaluation aiming to answer the following three questions empirically: (Q1) Section 5.1: *Does our method contribute to improve the generalization performance?* (Q2) Section 5.2: *What roles do pre-training phase and retraining phase play in our method?* (Q3) Section 5.3: *Can our method work on more complex models, like MetaOptNet [16] and CAVIA [36]?*

5.1 Few-Shot Classification Performances

We first evaluate the prediction performance of our method for few-shot classification tasks on two popular benchmark datasets: MiniImageNet [31] and TieredImageNet [24]. We have also evaluated our method on Omniglot [14] with numerical results relegated to Appendix C.1 due to space limit. The network used in our experiments is consistent with that considered for Reptile[21]. We test with varying channel number $\{32, 64, 128, 256\}$ in each convolution layer to show the robustness of our algorithms to meta-overfitting. See Appendix B for more details about Model, datasets and hyperparameters.

MiniImageNet The MiniImageNet dataset consists of 64 training classes, 12 validation classes and 24 test classes. For DSD-based Reptile, with 32 channels, we set the iteration numbers for the pre-training, pruning and retraining phases respectively as 3×10^4 , 5×10^4 and 2×10^4 , while with 64, 128, 256 channels, the corresponding number is 3×10^4 , 6×10^4 and 10^4 respectively. For IHT-based Reptile model training, we first pre-train the model for 2×10^4 iterations. Then we iterate between the sparse model fine-tuning (with 1.5×10^4 iterations) and dense-model retraining (with 5×10^3 iterations) for $t = 4$ rounds. The setting of other model training related parameters is identical to those in [21].

Table 1. Results on MiniImageNet under varying number of channels and pruning rates.

Methods	Backbone	Rate	5-way 1-shot	5-way 5-shot
Reptile baseline	32-32-32-32	-	50.30±0.40%	64.27±0.44%
	64-64-64-64	-	51.08±0.44%	65.46±0.43%
	128-128-128-128	-	49.96±0.45%	64.40±0.43%
	256-256-256-256	-	48.60±0.44%	63.24±0.43%
DSD+Reptile	32-32-32-32	40%	50.83±0.45%	65.24±0.44%
	64-64-64-64	30%	51.91±0.45%	67.23±0.43%
	128-128-128-128	50%	52.08±0.45%	68.87±0.42%
	256-256-256-256	60%	53.00±0.45%	68.04±0.42%
IHT+Reptile	32-32-32-32	20%	50.26±0.47%	63.63±0.45%
	64-64-64-64	40%	52.59±0.45%	67.41±0.43%
	128-128-128-128	40%	52.73±0.45%	68.69±0.42%
	256-256-256-256	60%	49.85±0.44%	66.56±0.42%

Results. The experimental results are presented in Table 1 and some additional results are provided in Table 4 in Appendix C. From these results, we can observe that our methods consistently outperform the considered baselines.

In the 32-channel setting in which the model is less prone to overfit, when applying DSD-based Reptile with 40% pruning rate, the accuracy gain is 0.5% on 5-way 1-shot tasks and 1% on 5-way 5-shot tasks. In the 64-channel setting, our IHT-based Reptile approach respectively improves about 1.5% and

1.95% over the baselines on 5-way 1-shot tasks and 5-way 5-shot tasks. In the setting of 128-channel, the accuracy of DSD-based Reptile on 5-way 1-shot tasks is nearly 3% higher than the baseline while on 5-way 5-shot tasks the gain is about 4.47%. Although in 256-channel case the performance of IHT-based approach drops compared with the 128-channel setting, it still achieves $\sim 1.2\%$ accuracy gain over the baseline on 5-way 1-shot tasks and $\sim 3.32\%$ on 5-way 5-shot tasks. These results clearly confirm the robustness of our algorithms to the meta-overfitting suffered from the over-parameterization of CNNs.

Figure 1 shows the evolving curves of training and testing accuracy under varying pruning rates from 0% to 40% for DSD and IHT based Reptile. From these curves we can clearly observe that the gap between training accuracy and testing accuracy reduces when the pruning rate increases, which confirms the predictions of Theorem 1 about the impact of network capacity on generalization.

Moreover, according to our knowledge, CAVIA [36], which introduces additional context parameters in the model, is also an effective approach for alleviating overfitting. In our experiment, we also compare our method with CAVIA. As shown in Table 4 in Appendix C.2, We can observe that our method can outperform CAVIA in most cases.




Fig. 1. The generalization performance of DSD-based and IHT-based Reptile on 5-way 1-shot and 5-way 5-shot tasks under 64-channel settings. **Left:** DSD-based Reptile; **Right:** IHT-based Reptile.

TieredImageNet The TieredImageNet dataset consists of 351 training classes, 97 validation classes and 160 test classes. For TieredImageNet dataset [24], in DSD-based Reptile case, we set iteration numbers for the pre-training, pruning and retraining phase respectively as 3×10^4 , 5×10^4 and 2×10^4 for all cases. In IHT-based Reptile, the values of iteration number are the same as those used in the previous experiments for MiniImageNet.

Results. The experimental results are partly presented in Table 2. More experimental results are available in Table 5 in Appendix C. In 5-way 1-shot classification tasks, both DSD-based Reptile approach and IHT-based Reptile approach outperform the baselines in all cases. In 32-channel setting, with DSD-based Reptile approach, the improvement of accuracy is 0.42% compared with baseline. In 64-channel setting, the accuracies of DSD-based Reptile and IHT-based Reptile respectively achieve 0.64% and 1.24% improvements. And for 256-channel setting, the best performance is also 0.44% better than baseline.

Table 2. Results on TieredImageNet under varying number of channels and pruning rates.

Methods	Backbone	Rate	5-way 1-shot	5-way 5-shot
Reptile baseline	32-32-32-32	-	50.52±0.45%	64.63±0.44%
	64-64-64-64	-	51.98±0.45%	67.70±0.43%
	128-128-128-128	-	53.30±0.45%	69.29±0.42%
	256-256-256-256	-	54.62±0.45%	68.06±0.42%
DSD+Reptile	32-32-32-32	10%	50.94±0.46%	64.65±0.44%
	64-64-64-64	10%	52.62±0.46%	66.69±0.43%
	128-128-128-128	10%	53.39±0.46%	67.22±0.43%
	256-256-256-256	20%	54.98±0.45%	67.98±0.43%
IHT+Reptile	32-32-32-32	10%	50.58±0.46%	63.09±0.45%
	64-64-64-64	20%	53.22±0.46%	66.15±0.44%
	128-128-128-128	10%	53.48±0.45%	69.39±0.42%
	256-256-256-256	10%	55.06±0.45%	67.60±0.43%

However, in most 5-way 5-shot classification tasks, the performance of our method drops. We conjecture that the reason is TieredImageNet dataset, compared with MiniImageNet dataset, contains more classes.

5.2 On the Impact of Hyperparameters

We next conduct a set of experiments on MiniImageNet to better understand the impact of pre-training and dense retraining on the task-specific testing performance.

Table 3. Results of ablation study in the 5-way setting. The “±” shows 95% confidence intervals, the “P.T” means “Pre-training” and the “R.T” means “Retraining”.

Methods	P.T	R.T	5-way 1-shot	5-way 5-shot
Reptile baseline(64)	-	-	51.08±0.44%	65.46±0.43%
Reptile baseline(128)	-	-	49.96±0.45%	64.40±0.43%
DSD+Reptile(64, 40%)	✓	×	43.92±0.43%	60.09±0.45%
DSD+Reptile(128, 60%)	✓	×	47.06±0.44%	55.07±0.44%
IHT+Reptile(64, 40%)	✓	×	40.03±0.41%	60.59±0.45%
IHT+Reptile(128, 60%)	✓	×	42.01±0.42%	52.71±0.45%
DSD+Reptile(64, 40%)	×	✓	50.84±0.45%	66.32±0.44%
DSD+Reptile(128, 60%)	×	✓	51.04±0.45%	67.23±0.44%
IHT+Reptile(64, 40%)	×	✓	52.07±0.45%	66.90±0.43%
IHT+Reptile(128, 60%)	×	✓	52.58±0.45%	67.83±0.42%
DSD+Reptile(64, 40%)	✓	✓	51.96±0.45%	66.64±0.43%
DSD+Reptile(128, 60%)	✓	✓	52.27±0.45%	68.44±0.42%
IHT+Reptile(64, 40%)	✓	✓	52.59±0.45%	67.41±0.43%
IHT+Reptile(128, 60%)	✓	✓	52.95±0.45%	68.04±0.42%

We begin by performing ablation study on pre-training and retraining phases.

Impact of the Retraining phase. It can be clearly seen from group of results in Table 3 that the retraining phase plays an important role in the accuracy performance of our method. Under the same pruning rate, without the retraining phase, the accuracy of both DSD-based and IHT-based Reptile approach drops dramatically. For an instance, in the 64-channel case with 40% pruning rate, the variant of IHT-based Reptile without retraining phase suffers from a $\sim 11\%$ drop in accuracy compared with the baseline. On the other side, as shown in Figure 2 that sparsity structure of the network does help to reduce the gap between training accuracy and testing accuracy even without the retraining phase. This confirms the benefit of sparsity for generalization gap reduction as revealed by Theorem 1. Therefore, the retraining phase helps to restore the capacity of the model to further improve the overall generalization performance.

Impact of the Pre-training phase. From Table 3, we observe that without pre-training phase, the variant algorithms still outperform baselines. Such results demonstrate the importance of pruning and retraining phase from another perspective that merely pruning and retraining the over-parameterized models can achieve similar empirical performance to our method. However, the variant algorithms fail to outperform our method. Since in network pruning, pre-training phase is treated as a necessary phase used to find a set of model parameters which is important [6, 7], we conjecture that it is the prematurely pruning before the model being well trained that leads to the drop of the performance.

We now perform experiments to further show how performance varies with different hyperparameters. The tested hyperparameters include (1)The number of pre-training and retraining iterations in DSD-based Reptile; (2) the number of iterations in an IHT pruning-retraining interval; (3) the ratio of pruning iterations in an interval. To be clear, we define $ratio = (Iter_{prune}/Iter_{interval})\%$. In experiments above, we set $Iter_{prune} = 1.5 \times 10^4$ in a 20000-iteration IHT interval, which means the ratio is 75%.

DSD-based Reptile. Figure 3 manifests the performance of DSD-based Reptile varying with the two hyperparameters, number of pre-training and retraining iterations. Figure 3(a) reveals that for most cases, too much or too little pre-training will both lead to the deterioration of performance. This is consis-




Fig. 2. Ablation study on retraining phase for both DSD-based Reptile and IHT-based Reptile on 64-channel case. The gap between the training accuracy and test accuracy of the variant algorithm of our method becomes smaller than that of baseline.




Fig. 3. Study of hyperparameters of DSD-based Reptile. (a). Study on the number of pre-training iterations. (b). Study on the number of retraining iterations.




Fig. 4. Study on hyperparameters of IHT-based Reptile. (a). Study on ratio of pruning iterations. (b). Study on the number of interval iterations.

tent with our ablation study that pre-training helps find a set of robust sparse parameters that is important and excessive pre-training, which reduces the iterations of pruning phase, undermines the generalization performance. Figure 3(b) shows that better performance can be obtained when retraining iterations are smaller than 30K, which indicates that only a small number of retraining steps are required to restore the accuracy without overfitting again.

IHT-based Reptile. Figure 4 shows the hyper-parameter sensitivity results of IHT-based Reptile. Figure 4(a) shows the performance under different ratio of pruning iterations in an IHT interval. It’s clear that better performance can be obtained when the ratio is larger than 50%, which means pruning iterations are more than retraining iterations. This reveals that more pruning iterations are required to alleviate overfitting and a small number of retraining steps are enough to help compensate the loss of accuracy. Figure 4(b) shows the performance under varying number of iterations in an IHT interval. We can see that with the interval iterations increasing from 5K to 20K, the accuracies get improved. This suggests that sufficient steps are required to train a robust model in a loop of pruning-retraining.

5.3 Performance on More Complex Networks

We further implement our method on *MetaOptNet* [16] and *CAVIA* [36] to evaluate its the performance on more complex network structures. For MetaOpeNet,

Table 4. Results on complex Networks. The “ \pm ” shows 95% confidence intervals.

Methods	Backbone	Rate	5-way 1-shot	5-way 5-shot
MetaOptNet(rerun) [16]	ResNet-12	-	61.95 \pm 0.60%	77.79 \pm 0.45%
DSD-MetaOptNet	ResNet-12	20%	62.16\pm0.62%	77.51 \pm 0.48%
CAVIA(rerun) [36]	128-128-128-128	-	48.76 \pm 0.99%	62.54 \pm 0.78%
DSD-CAVIA	128-128-128-128	10%	49.53\pm0.93%	63.34\pm0.79%
IHT-CAVIA	128-128-128-128	10%	49.97\pm0.97%	63.48\pm0.78%

we select the ResNet-12 as the network and SVM as the head, and the dropout in the network is replaced by our method. For CAVIA, our methods are applied only on the network parameters. Other settings are the same as those in paper [36]. The dataset used is MiniImageNet. More experimental settings can be referred to Appendix B.3.

As shown in Table 4, for MetaOptNet, our method gains 0.2% improvement on 5-way 1-shot tasks compared with baseline. On 5-way 5-shot tasks, our method can still obtain similar performance. However, there is a trade-off between accuracy and training time. For CAVIA, all cases outperform the baselines, which shows the strong power of our methods in alleviating the overfitting.

6 Conclusion

In this paper, we proposed a cardinality-constrained meta-learning approach for improving generalization performance via explicitly controlling the capacity of over-parameterized neural networks. We have theoretically proved that the generalization gap bounds of the sparse meta-learner have polynomial dependence on the sparsity level rather than the number of parameters. Our approach has been implemented in a scalable meta-learning framework of Reptile with the sparsity level of parameters maintained by network pruning routines including dense-sparse-dense and iterative hard thresholding. Extensive experimental results on benchmark few-shot classification tasks, along with hyperparameter impact study and the study on complex networks, confirm our theoretical predictions and demonstrate the power of network pruning and retraining for improving the generalization performance of gradient-optimization-based meta-learning.

Acknowledgements

Xiao-Tong Yuan is supported in part by National Major Project of China for New Generation of AI under Grant No.2018AAA0100400 and in part by Natural Science Foundation of China (NSFC) under Grant No.61876090 and No.61936005. Qingshan Liu is supported by NSFC under Grant No.61532009 and No.61825601.

References

1. Abramovich, F., Grinshtein, V.: High-dimensional classification by sparse logistic regression. *IEEE Transactions on Information Theory* **65**(5), 3068–3079 (2019)
2. Arora, S., Ge, R., Neyshabur, B., Zhang, Y.: Stronger generalization bounds for deep nets via a compression approach. In: *International Conference on Machine Learning*. pp. 254–263 (2018)
3. Bengio, Y., Bengio, S., Cloutier, J.: Learning a synaptic learning rule. In: *IJCNN* (1990)
4. Denevi, G., Ciliberto, C., Grazi, R., Pontil, M.: Learning-to-learn stochastic gradient descent with biased regularization. In: *International Conference on Machine Learning*. pp. 1566–1575 (2019)
5. 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)
6. Frankle, J., Carbin, M.: The lottery ticket hypothesis: Finding sparse, trainable neural networks. In: *International Conference on Learning Representations* (2019)
7. Han, S., Pool, J., Narang, S., Mao, H., Gong, E., Tang, S., Elsen, E., Vajda, P., Paluri, M., Tran, J., et al.: Dsd: Dense-sparse-dense training for deep neural networks. In: *International Conference on Learning Representations* (2016)
8. Han, S., Pool, J., Tran, J., Dally, W.: Learning both weights and connections for efficient neural network. In: *Advances in Neural Information Processing Systems*. pp. 1135–1143 (2015)
9. Hardt, M., Recht, B., Singer, Y.: Train faster, generalize better: Stability of stochastic gradient descent. In: *International Conference on Machine Learning*. pp. 1225–1234 (2016)
10. Hassibi, B., Stork, D.G., Wolff, G.J.: Optimal brain surgeon and general network pruning. In: *IEEE International Conference on Neural Networks*. pp. 293–299. *IEEE* (1993)
11. Jin, X., Yuan, X., Feng, J., Yan, S.: Training skinny deep neural networks with iterative hard thresholding methods. *arXiv preprint arXiv:1607.05423* (2016)
12. Khodak, M., Balcan, M.F., Talwalkar, A.: Provable guarantees for gradient-based meta-learning. In: *Advances in Neural Information Processing Systems* (2019)
13. Koch, G., Zemel, R., Salakhutdinov, R.: Siamese neural networks for one-shot image recognition. In: *ICML deep learning workshop*. vol. 2 (2015)
14. Lake, B., Salakhutdinov, R., Gross, J., Tenenbaum, J.: One shot learning of simple visual concepts. In: *Proceedings of the annual meeting of the cognitive science society*. vol. 33 (2011)
15. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: *Advances in Neural Information Processing Systems*. pp. 598–605 (1990)
16. 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)
17. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few-shot learning. In: *Advances in Neural Information Processing Systems* (2017)
18. Louizos, C., Welling, M., Kingma, D.P.: Learning sparse neural networks through l_0 regularization. *arXiv preprint arXiv:1712.01312* (2017)
19. Maurer, A., Pontil, M.: Structured sparsity and generalization. *Journal of Machine Learning Research* **13**(Mar), 671–690 (2012)

20. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=B1DmUzWAW>
21. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999 (2018)
22. Pillutla, V.K., Roulet, V., Kakade, S.M., Harchaoui, Z.: A smoother way to train structured prediction models. In: Advances in Neural Information Processing Systems. pp. 4766–4778 (2018)
23. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: International Conference on Learning Representations (2016)
24. Ren, M., Triantafillou, E., Ravi, S., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. arXiv preprint arXiv:1803.00676 (2018)
25. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-learning with memory-augmented neural networks. In: International Conference on Machine Learning. pp. 1842–1850 (2016)
26. Schmidhuber, J.: Evolutionary principles in self-referential learning. On learning how to learn: The meta-meta-... hook.) Diploma thesis, Institut f. Informatik, Tech. Univ. Munich **1**, 2 (1987)
27. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: Advances in Neural Information Processing Systems. pp. 4077–4087 (2017)
28. Srinivas, S., Babu, R.V.: Data-free parameter pruning for deep neural networks. arXiv preprint arXiv:1507.06149 (2015)
29. Sung, F., Zhang, L., Xiang, T., Hospedales, T., Yang, Y.: Learning to learn: Meta-critic networks for sample efficient learning. arXiv preprint arXiv:1706.09529 (2017)
30. Thrun, S., Pratt, L.: Learning to learn. Springer Science & Business Media (2012)
31. 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)
32. Weston, J., Chopra, S., Bordes, A.: Memory networks. arXiv preprint arXiv:1410.3916 (2014)
33. Yoon, J., Kim, T., Dia, O., Kim, S., Bengio, Y., Ahn, S.: Bayesian model-agnostic meta-learning. In: Advances in Neural Information Processing Systems. pp. 7332–7342 (2018)
34. Yuan, X.T., Li, P., Zhang, T.: Gradient hard thresholding pursuit. *Journal of Machine Learning Research* **18**, 1–43 (2018)
35. Zhou, P., Yuan, X., Xu, H., Yan, S., Feng, J.: Efficient meta learning via mini-batch proximal update. In: Advances in Neural Information Processing Systems. pp. 1532–1542 (2019)
36. Zintgraf, L., Shiarli, K., Kurin, V., Hofmann, K., Whiteson, S.: Fast context adaptation via meta-learning. In: International Conference on Machine Learning. pp. 7693–7702 (2019)