



Regularizing Meta-learning via Gradient Dropout

Hung-Yu Tseng¹(✉), Yi-Wen Chen¹, Yi-Hsuan Tsai², Sifei Liu³, Yen-Yu Lin⁴,
and Ming-Hsuan Yang^{1,5}

¹ University of California, Merced, Merced, USA
{[htseng6](mailto:htseng6@ucmerced.edu), [ychen319](mailto:ychen319@ucmerced.edu)}@ucmerced.edu

² NEC Laboratories America, Princeton, USA

³ Nvidia Research, Santa Clara, USA

⁴ National Chiao Tung University, Hsinchu, Taiwan

⁵ Google Research, Cambridge, USA

Abstract. With the growing attention on learning-to-learn new tasks using only a few examples, meta-learning has been widely used in numerous problems such as few-shot classification, reinforcement learning, and domain generalization. However, meta-learning models are prone to overfitting when there are no sufficient training tasks for the meta-learners to generalize. Although existing approaches such as Dropout are widely used to address the overfitting problem, these methods are typically designed for regularizing models of a single task in supervised training. In this paper, we introduce a simple yet effective method to alleviate the risk of overfitting for gradient-based meta-learning. Specifically, during the gradient-based adaptation stage, we randomly drop the gradient in the inner-loop optimization of each parameter in deep neural networks, such that the augmented gradients improve generalization to new tasks. We present a general form of the proposed gradient dropout regularization and show that this term can be sampled from either the Bernoulli or Gaussian distribution. To validate the proposed method, we conduct extensive experiments and analysis on numerous computer vision tasks, demonstrating that the gradient dropout regularization mitigates the overfitting problem and improves the performance upon various gradient-based meta-learning frameworks.

1 Introduction

In recent years, significant progress has been made in meta-learning, which is also known as *learning to learn*. One common setting is that, given only a few training examples, meta-learning aims to learn new *tasks* rapidly by leveraging

H.-Y. Tseng and Y.-W. Chen—Equal contribution.

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-69538-5_14) contains supplementary material, which is available to authorized users.

© Springer Nature Switzerland AG 2021

H. Ishikawa et al. (Eds.): ACCV 2020, LNCS 12625, pp. 218–234, 2021.

https://doi.org/10.1007/978-3-030-69538-5_14

the past experience acquired from the known tasks. It is a vital machine learning problem due to the potential for reducing the amount of data and time for adapting an existing system. Numerous recent methods successfully demonstrate how to adopt meta-learning algorithms to solve various learning problems, such as few-shot classification [1–3], reinforcement learning [4, 5], and domain generalization [6, 7].

Despite the demonstrated success, meta-learning frameworks are prone to overfitting [8] when there do not exist sufficient training tasks for the meta-learners to generalize. For instance, the mini-ImageNet [9] few-shot classification dataset contains only 64 training categories. Since the training tasks can be only sampled from this small set of classes, meta-learning models may overfit and fail to generalize to new testing tasks.

Significant efforts have been made to address the overfitting issue in the supervised learning framework, where the model is developed to learn a *single* task (e.g., recognizing the same set of categories in both training and testing phases). The Dropout [10] method randomly drops (zeros) intermediate activations in deep neural networks during the training stage. Relaxing the limitation of binary dropout, the Gaussian dropout [11] scheme augments activations with noise sampled from a Gaussian distribution. Numerous methods [12–16] further improve the Dropout method by injecting structural noise or scheduling the dropout process to facilitate the training procedure. Nevertheless, these methods are developed to regularize the models to learn a single task, which may not be effective for meta-learning frameworks.

In this paper, we address the overfitting issue [8] in gradient-based meta-learning. As shown in Fig. 1(a), given a new task, the meta-learning framework aims to adapt model parameters θ to be θ' via the gradients computed according to the few examples (support data \mathcal{X}^s). This gradient-based adaptation process is also known as the *inner-loop* optimization. To alleviate the overfitting issue, one straightforward approach is to apply the existing dropout method to the model weights directly. However, there are two sets of model parameters θ and θ' in the inner-loop optimization. As such, during the meta-training stage, applying normal dropout would cause inconsistent randomness, i.e., dropped neurons, between these two sets of model parameters. To tackle this issue, we propose a dropout method on the gradients in the inner-loop optimization, denoted as *DropGrad*, to regularize the training procedure. This approach naturally bridges θ and θ' , and thereby involves only one randomness for the dropout regularization. We also note that our method is model-agnostic and generalized to various gradient-based meta-learning frameworks such as [1, 17, 18]. In addition, we demonstrate that the proposed dropout term can be formulated in a general form, where either the binary or Gaussian distribution can be utilized to sample the noise, as demonstrated in Fig. 1(b).

To evaluate the proposed DropGrad method, we conduct experiments on numerous computer vision tasks, including few-shot classification on the mini-ImageNet [9], online object tracking [19], and few-shot viewpoint estimation [20], showing that the DropGrad scheme can be applied to and improve different tasks. In addition, we present comprehensive analysis by using various meta-learning frameworks, adopting different dropout probabilities, and explaining

which layers to apply gradient dropout. To further demonstrate the generalization ability of DropGrad, we perform a challenging cross-domain few-shot classification task, in which the meta-training and meta-testing sets are from two different distributions, *i.e.*, the mini-ImageNet and CUB [21] datasets. We show that with the proposed method, the performance is significantly improved under the cross-domain setting. Our source code is available at <https://github.com/hytseng0509/DropGrad>.

In this paper, we make the following contributions:

- We propose a simple yet effective gradient dropout approach to improve the generalization ability of gradient-based meta-learning frameworks.
- We present a general form for gradient dropout and show that both binary and Gaussian sampling schemes mitigate the overfitting issue.
- We demonstrate the effectiveness and generalizability of the proposed method via extensive experiments on numerous computer vision tasks.

2 Related Work

Meta-Learning. Meta-learning aims to adapt the past knowledge learned from previous tasks to new tasks with few training instances. Most meta-learning algorithms can be categorized into three groups: 1) Memory-based approaches [2, 22] utilize recurrent networks to process few training examples of new tasks sequentially; 2) Metric-based frameworks [3, 9, 23, 24] make predictions by referring to the features encoded from the input data and training instances in a generic metric space; 3) Gradient-based methods [1, 8, 17, 18, 25–27] learn to optimize the model via gradient descent with few examples, which is the focus of this work. In the third group, the MAML [1] approach learns model initialization (*i.e.*, initial parameters) that is amenable to fast fine-tuning with few instances. In addition to model initialization, the MetaSGD [18] method learns a set of learning rates for different model parameters. Furthermore, the MAML++ [17] algorithm makes several improvements based on the MAML method to facilitate the training process with additional performance gain. However, these methods are still prone to overfitting as the dataset for the training tasks is insufficient for the model to adapt well. Recently, Kim *et al.* [8] and Rusu *et al.* [26] address this issue via the Bayesian approach and latent embeddings. Nevertheless, these methods employ additional parameters or networks which entail significant computational overhead and may not be applicable to arbitrary frameworks. In contrast, the proposed gradient dropout regularization does not impose any overhead and thus can be readily integrated into the gradient-based models mentioned above.

Dropout Regularization. Built upon the Dropout [10] method, various schemes [12–15, 28] have been proposed to regularize the training process of deep neural networks for supervised learning. The core idea is to inject noise into intermediate activations when training deep neural networks. Several recent studies improve the regularization on convolutional neural networks by making the injected structural noise. For instance, the SpatialDropout [14] method drops

the entire channel from an activation map, the DropPath [13, 16] scheme chooses to discard an entire layer, and the DropBlock [12] algorithm zeros multiple continuous regions in an activation map. Nevertheless, these approaches are designed for deep neural networks that aim to learn a *single* task, *e.g.*, learning to recognize a fixed set of categories. In contrast, our algorithm aims to regularize the gradient-based meta-learning frameworks that suffer from the overfitting issue on the *task*-level, *e.g.*, introducing new tasks.




Fig. 1. Illustration of the proposed method. (a) The proposed DropGrad method imposes a noise term n to augment the gradient in the inner-loop optimization during the meta-training stage. (b) The DropGrad method samples the noise term n from either the Bernoulli or Gaussian distribution, in which the Gaussian distribution provides a better way to account for uncertainty.

3 Gradient Dropout Regularization

Before introducing details of our proposed dropout regularization on gradients, we first review the gradient-based meta-learning framework.

3.1 Preliminaries for Meta-learning

In meta-learning, multiple tasks $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$ are divided into meta-training $\mathcal{T}^{\text{train}}$, meta-validation \mathcal{T}^{val} , and meta-testing $\mathcal{T}^{\text{test}}$ sets. Each task T_i consists of a support set $D^s = (\mathcal{X}^s, \mathcal{Y}^s)$ and a query set $D^q = (\mathcal{X}^q, \mathcal{Y}^q)$, where \mathcal{X} and \mathcal{Y} are a set of input data and the corresponding ground-truth. The support set D^s represents the set of few labeled data for learning, while the query set D^q indicates the set of data to be predicted.

Given a novel task and a parametric model f_θ , the objective of a gradient-based approach during the meta-training stage is to minimize the prediction loss L^q on the query set D^q according to the signals provided from the support set D^s , and thus the model f_θ can be adapted. Figure 1(a) shows an overview of the MAML [1] method, which offers a general formulation of gradient-based frameworks. For each iteration of the meta-training phase, we first randomly

sample a task $T = \{D^s, D^q\}$ from the meta-training set $\mathcal{T}^{\text{train}}$. We then adapt the initial parameters θ to be task-specific parameters θ' via gradient descent:

$$\theta' = \theta - \alpha \odot g, \quad (1)$$

where α is the learning rate for gradient-based adaptation and \odot is the operation of element-wise product, *i.e.*, Hadamard product. The term g in (1) is the set of gradients computed according to the objectives of model f_θ on the support set $D^s = (\mathcal{X}^s, \mathcal{Y}^s)$:

$$g = \nabla_\theta L^s(f_\theta(\mathcal{X}^s), \mathcal{Y}^s). \quad (2)$$

We call the step of (1) as the inner-loop optimization and typically, we can do multiple gradient steps for (1), *e.g.*, smaller than 10 in general. After the gradient-based adaptation, the initial parameters θ are optimized according to the loss functions of the adapted model $f_{\theta'}$ on the query set $D^q = (\mathcal{X}^q, \mathcal{Y}^q)$:

$$\theta = \theta - \eta \nabla_\theta L^q(f_{\theta'}(\mathcal{X}^q), \mathcal{Y}^q), \quad (3)$$

where η is the learning rate for meta-training. During the meta-testing stage, the model f_θ is adapted according to the support set D^s and the prediction on query data \mathcal{X}^q is made without accessing the ground-truth \mathcal{Y}^q in the query set. We note that several methods are built upon the above formulation introduced in the MAML method. For example, the learning rate α for gradient-adaptation is viewed as the optimization objective [17,18], and the initial parameters θ are not generic but conditional on the support set D^s [26].

3.2 Gradient Dropout

The main idea is to impose uncertainty to the core objective during the meta-training step, *i.e.*, the gradient g in the inner-loop optimization, such that θ' receives gradients with noise to improve the generalization of gradient-based models. As described in Sect. 3.1, adapting the model θ to θ' involves the gradient update in the inner-loop optimization formulated in (2). Based on this observation, we propose to randomly drop the gradient in (2), *i.e.*, g , during the inner-loop optimization, as illustrated in Fig. 1. Specifically, we augment the gradient g as follows:

$$g' = g \odot n, \quad (4)$$

where n is a noise regularization term sampled from a pre-defined distribution. With the formulation of (4), in the following we introduce two noise regularization strategies via sampling from different distributions, *i.e.*, the Bernoulli and Gaussian distributions.

Binary DropGrad. We randomly zero the gradient with the probability p , in which the process can be formulated as:

$$g' = g \odot n_b, \quad n_b \sim \frac{\text{Bernoulli}(1-p)}{1-p}, \quad (5)$$

where the denominator $1-p$ is the normalization factor. Note that, different from the Dropout [10] method which randomly drops the intermediate activations in a supervised learning network under a single task setting, we perform the dropout on the gradient level.

Gaussian DropGrad. One limitation of the Binary DropGrad scheme is that the noise term n_b is only applied in a binary form, which is either 0 or $1-p$. To address this disadvantage and provide a better regularization with uncertainty, we extend the Bernoulli distribution to the Gaussian formulation. Since the expectation and variance of the noise term n_b in the Binary DropGrad method are respectively $E(n_b) = 1$ and $\sigma^2(n_b) = \frac{p}{1-p}$, we can augment the gradient g with noise sampled from the Gaussian distribution:

$$g' = g \odot n_g, \quad n_g \sim \text{Gaussian}(1, \frac{p}{1-p}). \quad (6)$$

Algorithm 1: Applying DropGrad on MAML [1]

```

1 Require: a set of training tasks  $\mathcal{T}^{\text{train}}$ , adaptation learning rate  $\alpha$ ,
   meta-learning rate  $\eta$ 
2 randomly initialize  $\theta$ 
3 while training do
4   randomly sample a task  $T = \{D^s(\mathcal{X}^s, \mathcal{Y}^s), D^q(\mathcal{X}^q, \mathcal{Y}^q)\}$  from  $\mathcal{T}^{\text{train}}$ 
5    $g = \nabla_{\theta} L^s(f_{\theta}(\mathcal{X}^s), \mathcal{Y}^s)$ 
6   compute  $g'$  according to (5) or (6) // Apply DropGrad
7    $\theta' = \theta - \alpha \times g'$ 
8    $\theta = \theta - \eta \nabla_{\theta} L^q(f_{\theta'}(\mathcal{X}^q), \mathcal{Y}^q)$ 
9 end

```

As a result, two noise terms n_b and n_g are statistically comparable with the same dropout probability p . In Fig. 1(b), we illustrate the difference between the Binary DropGrad and Gaussian DropGrad approaches. We also show the process of applying the proposed regularization using the MAML [1] method in Algorithm 1, while similar procedures can be applied to other gradient-based meta-learning frameworks, such as MetaSGD [18] and MAML++ [17].

4 Experimental Results

In this section, we evaluate the effectiveness of the proposed DropGrad method by conducting extensive experiments on three learning problems: few-shot classification, online object tracking, and few-shot viewpoint estimation. In addition, for the few-shot classification experiments, we analyze the effect of using binary and Gaussian noise, which layers to apply DropGrad, and performance in the cross-domain setting.

4.1 Few-Shot Classification

Few-shot classification aims to recognize a set of new categories, *e.g.*, five categories (5-way classification), with few, *e.g.*, one (1-shot) or five (5-shot), example images from each category. In this setting, the support set D^s contains the few images \mathcal{X}^s of the new categories and the corresponding categorical annotation \mathcal{Y}^s . We conduct experiments on the mini-ImageNet [9] dataset, which is widely used for evaluating few-shot classification approaches. As a subset of the ImageNet [29], the mini-ImageNet dataset contains 100 categories and 600 images for each category. We use the 5-way evaluation protocol in [30] and split the dataset into 64 training, 16 validating, and 20 testing categories.




Fig. 2. Comparison between the proposed Binary and Gaussian DropGrad methods. We compare the 1-shot (*left*) and 5-shot (*right*) performance of MAML [1] trained with two different forms of DropGrad under various dropout rates on mini-ImageNet. The proposed DropGrad method is particularly effective with the dropout rate in [0.1, 0.2]. Moreover, the Gaussian DropGrad method consistently obtains better results compared to the Binary DropGrad scheme. Therefore, we apply the Gaussian DropGrad method with the dropout rate of 0.1 or 0.2 in all of our experiments.

Implementation Details. We apply the proposed DropGrad regularization method to train the following gradient-based meta-learning frameworks: MAML [1], MetaSGD [18], and MAML++ [17]. We use the implementation from Chen *et al.* [31] for MAML and use our own implementation for MetaSGD.¹ We use the ResNet-18 [32] model as the backbone network for both MAML and MetaSGD. As for MAML++, we use the original source code.² Similar to recent studies [26], we also pre-train the feature extractor of ResNet-18 by minimizing the classification loss on the 64 training categories from the mini-ImageNet dataset for the MetaSGD method, which is denoted by MetaSGD*.

For all the experiments, we use the *default hyper-parameter settings* provided by the original implementation. Moreover, we select the model according to the validation performance for evaluation (*i.e.*, *early stopping strategy*).

Comparison between Binary and Gaussian DropGrad. We first evaluate how the proposed Binary and Gaussian DropGrad methods perform on the

¹ <https://github.com/wyharveychen/CloserLookFewShot>.

² <https://github.com/AntreasAntoniou/HowToTrainYourMAMLPytorch>.

MAML framework with different values of the dropout probability p . Figure 2 shows that both methods are effective especially when the dropout rate is in the range of $[0.1, 0.2]$, while setting the dropout rate to 0 is to turn the proposed DropGrad method off. Since the problem of learning from only one instance (1-shot) is more complicated, the overfitting effect is less severe compared to the 5-shot setting. As a result, applying the DropGrad method with a dropout rate larger than 0.3 degrades the performance. Moreover, the Gaussian DropGrad method consistently outperforms the binary case on both 1-shot and 5-shot tasks, due to a better regularization term n_g with uncertainty. We then apply the Gaussian DropGrad method with the dropout rate of 0.1 or 0.2 in the following experiments.

Table 1. Few-shot classification results on mini-ImageNet. The Gaussian DropGrad method improves the performance of gradient-based models on 1-shot and 5-shot classification tasks.

Model	1-shot	5-shot
MAML [1]	49.61 \pm 0.92%	65.72 \pm 0.77%
MAML w/ Gaussian DropGrad	52.35 \pm 0.86%	69.42 \pm 0.73%
MetaSGD [18]	51.51 \pm 0.87%	69.67 \pm 0.75%
MetaSGD w/ Gaussian DropGrad	53.38 \pm 0.93%	71.14 \pm 0.72%
MetaSGD*	60.44 \pm 0.87%	72.55 \pm 0.54%
MetaSGD* w/ Gaussian DropGrad	61.69 \pm 0.84%	73.33 \pm 0.57%
MAML++ [17]	50.21 \pm 0.50%	68.66 \pm 0.46%
MAML++ w/ Gaussian DropGrad	51.13 \pm 0.50%	69.80 \pm 0.46%

Table 2. Performance of applying DropGrad to different layers. We conduct experiments on the 5-shot classification task using MAML on mini-ImageNet. It is more helpful in improving the performance by dropping the gradients closer to the output layers (*e.g.*, FC and Block4 + FC).

Origin	FC	Block4 + FC	Full	Block1 + Conv	Conv
65.72 \pm 0.77%	68.93 \pm 0.55%	69.02 \pm 0.57%	69.42 \pm 0.73%	64.96 \pm 0.80%	65.53 \pm 0.75%

Comparison with Existing Dropout Methods. To show that the proposed DropGrad method is effective for gradient-based meta-learning frameworks, we compare it with two existing dropout schemes applied on the network activations in both f_θ and f'_θ . We choose the Dropout [10] and SpatialDropout [14] methods, since the former is a commonly-used approach while the latter is shown to be effective for applying to 2D convolutional maps. The performance of MAML on 5-shot classification on the mini-ImageNet dataset is: *DropGrad* 69.42 \pm 0.73%, *SpatialDropout* 68.09 \pm 0.56%, and *Vanilla Dropout* 67.44 \pm 0.57%. This demonstrates the benefit of using the proposed DropGrad method, which effectively

tackles the issue of inconsistent randomness between two different models f_θ and f'_θ in the inner-loop optimization of gradient-based meta-learning frameworks.

Overall Performance on the Mini-ImageNet Dataset. Table 1 shows the results of applying the proposed Gaussian DropGrad method to different frameworks. The results validate that the proposed regularization scheme consistently improves the performance of various gradient-based meta-learning approaches. In addition, we present the curve of validation loss over training episodes from MAML and MetaSGD on the 5-shot classification task in Fig. 3. We observe that the overfitting problem is more severe in training the MetaSGD method since it consists of more parameters to be optimized. The DropGrad regularization method mitigates the overfitting issue and facilitates the training procedure.

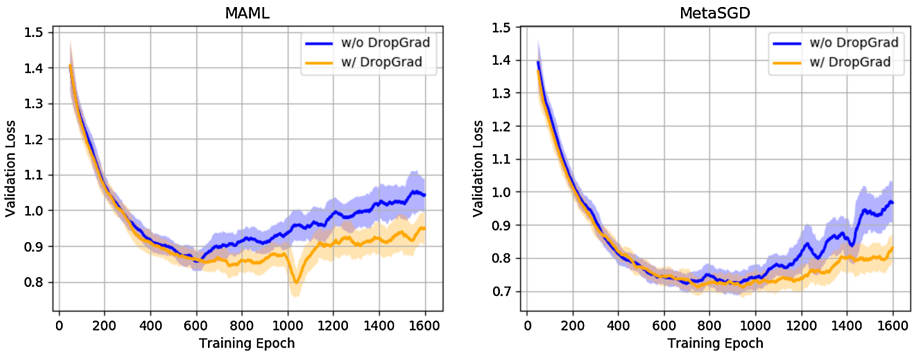


Fig. 3. Validation loss over training epochs. We show the validation curves of the MAML (*left*) and MetaSGD (*right*) frameworks trained on the 5-shot mini-ImageNet dataset. The curves and shaded regions represent the mean and standard deviation of validation loss over 50 epochs. The curves validate that the proposed DropGrad method alleviates the overfitting problem.

Table 3. 5-shot classification results of MAML under various hyperparameter settings. We study the learning rate α and number of iterations n_{inner} in the inner-loop optimization of MAML using mini-ImageNet dataset.

α, n_{inner}	0.01, 5 (original)	0.1, 5	0.001, 5	0.01, 3	0.01, 7
MAML [1]	65.72 \pm 0.77%	65.98 \pm 0.79%	58.55 \pm 0.80%	64.84 \pm 0.80%	68.11 \pm 0.74%
MAML w/ DropGrad	69.42 \pm 0.73%	67.78 \pm 0.73%	64.05 \pm 0.79%	65.42 \pm 0.80%	69.65 \pm 0.70%

Layers to Apply DropGrad. We study which layers in the network to apply the DropGrad regularization in this experiment. The backbone ResNet-18 model contains a convolutional layer (Conv) followed by 4 residual blocks (Block1, Block2, Block3, Block4) and a fully-connected layer (FC) as the classifier. We perform the Gaussian DropGrad method on different parts of the ResNet-18 model for MAML on the 5-shot classification task. The results are presented in

Table 2. We find that it is more critical to drop the gradients closer to the output layers (*e.g.*, FC and Block4 + FC). Applying the DropGrad method to the input side (*e.g.*, Block1 + Conv and Conv), however, may even negatively affect the training and degrade the performance. This can be explained by the fact that features closer to the output side are more abstract and thus tend to overfit. As using the DropGrad regularization term only increases a negligible overhead, we use the *Full* model, where our method is applied to all layers in the experiments unless otherwise mentioned.

Hyper-Parameter Analysis. In all experiments shown in Sect. 4, we use the default hyper-parameter values from the original implementation of the adopted methods. In this experiment, we explore the hyper-parameter choices for MAML [1]. Specifically, we conduct an ablation study on the learning rate α and the number of inner-loop optimizations n_{inner} in MAML. As shown in Table 3, the proposed DropGrad method improves the performance consistently under different sets of hyper-parameters.

Table 4. Cross-Domain performance for few-shot classification. We use the mini-ImageNet and CUB datasets for the meta-training and meta-testing steps, respectively. The improvement of applying the proposed DropGrad method is more significant in the cross-domain cases than the intra-domain ones.

Model	1-Shot	5-Shot
MAML [1]	31.52 \pm 0.52%	45.56 \pm 0.51%
MAML w/ Dropout [10]	31.84 \pm 0.49%	46.48 \pm 0.50%
MAML w/ DropGrad	33.20 \pm 0.67%	51.05 \pm 0.56%
MetaSGD [18]	34.52 \pm 0.63%	49.22 \pm 0.58%
MetaSGD w/ Dropout [10]	35.01 \pm 0.54%	52.35 \pm 0.58%
MetaSGD w/ DropGrad	36.77 \pm 0.72%	55.13 \pm 0.72%
MetaSGD*	43.98 \pm 0.77%	57.95 \pm 0.81%
MetaSGD* w/ DropGrad	45.33 \pm 0.81%	59.94 \pm 0.82%
MAML++ [17]	40.73 \pm 0.49%	60.57 \pm 0.49%
MAML++ w/ Dropout [10]	41.75 \pm 0.49%	61.48 \pm 0.49%
MAML++ w/ DropGrad	44.27 \pm 0.50%	63.79 \pm 0.48%

4.2 Cross-domain Few-Shot Classification

To further evaluate how the proposed DropGrad method improves the generalization ability of gradient-based meta-learning models, we conduct a cross-domain experiment, in which the meta-testing set is from an *unseen* domain. We use the cross-domain scenario introduced by Chen *et al.* [31], where the meta-training step is performed on the mini-ImageNet [9] dataset while the meta-testing evaluation is conducted on the CUB [33] dataset. Note that, different

from Chen *et al.* [31] who select the model according to the validation performance on the CUB dataset, we pick the model via the validation performance on the mini-ImageNet dataset for evaluation. The reason is that we target at

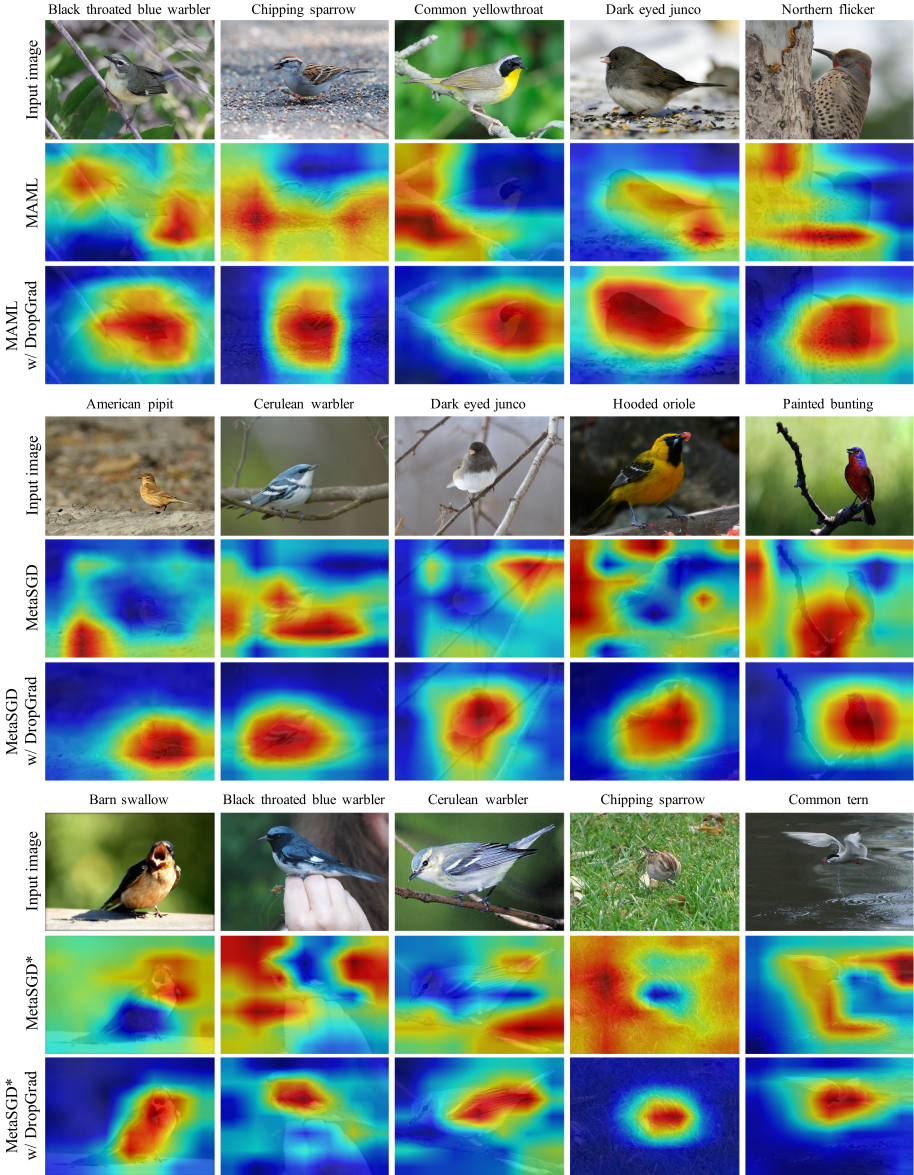


Fig. 4. Class activation maps (CAMs) for cross-domain 5-shot classification. The mini-ImageNet and CUB datasets are used for the meta-training and meta-testing steps, respectively. Models trained with the proposed DropGrad (the third row for each example) focus more on the objects than the original models (the second row for each example).

analyzing the generalization ability to the unseen domain, and thus we do not utilize any information provided from the CUB dataset.

Table 4 shows the results using the Gaussian DropGrad method. Since the domain shift in the cross-domain scenario is larger than that in the intra-domain case (*i.e.*, both training and testing tasks are sampled from the mini-ImageNet dataset), the performance gains of applying the proposed DropGrad method reported in Table 4 are more significant than those in Table 1. The results demonstrate that the DropGrad scheme is able to effectively regularize the gradients and transfer them for learning new tasks in an unseen domain.

Table 5. Precision and success rate on the OTB2015 dataset. The DropGrad method can be applied to visual tracking and improve the tracking performance.

Model	Precision	Success rate
MetaCREST [19]	0.7994	0.6029
MetaCREST w/ DropGrad	0.8172	0.6145
MetaSDNet [19]	0.8673	0.6434
MetaSDNet w/ DropGrad	0.8746	0.6520

To further understand the improvement by the proposed method under the cross-domain setting, we visualize the class activation maps (CAMs) [34] of the images in the unseen domain (CUB). More specifically, during the testing time, we adapt the learner model f_θ with the support set D^s . We then compute the class activation maps of the data in the query set D^q from the last convolutional layer of the updated learner model f'_θ . Figure 4 demonstrates the results of the MAML, MetaSGD, and MetaSGD* approaches. The models trained with the proposed regularization method show the activation on more discriminative regions. This suggests that the proposed regularization improves the generalization ability of gradient-based schemes, and thus enables these methods to adapt to the novel tasks sampled from the unseen domain.

Comparison with the Existing Dropout Approach. We also compare the proposed DropGrad approach with existing Dropout [10] method under the cross-domain setting. We apply the existing Dropout scheme on the network activations in both f_θ and f'_θ . As suggested by Ghiasi *et al.* [12], we use the dropout rate of 0.3 for the Dropout method. As the results shown in Table 4, the proposed DropGrad method performs favorably against the Dropout approach. The larger performance gain from the DropGrad approach validates effectiveness of imposing uncertainty on the inner-loop gradient for the gradient-based meta-learning framework. On the other hand, since applying the conventional Dropout causes the inconsistent randomnesses between two different sets of parameters f_θ and f'_θ , which is less effective compared to the proposed scheme.

4.3 Online Object Tracking

Visual object tracking targets at localizing one particular object in a video sequence given the bounding box annotation in the first frame. To adapt the model to the subsequent frames, one approach is to apply online adaptation during tracking. The Meta-Tracker [19] method uses meta-learning to improve two state-of-the-art online trackers, including the correlation-based CREST [35] and the detection-based MDNet [36], which are denoted as MetaCREST and MetaSDNet. Based on the error signals from future frames, the Meta-Tracker updates the model during offline meta-training, and obtains a robust initial network that generalizes well over future frames. We apply the proposed DropGrad method to train the MetaCREST and MetaSDNet models with evaluation on the OTB2015 [37] dataset.




Fig. 5. Qualitative results of object online tracking on the OTB2015 dataset. Red boxes are the ground-truth, yellow boxes represent the original results, and green boxes stand for the results where the DropGrad method is applied. Models trained with the proposed DropGrad scheme are able to track objects more accurately. (Color figure online)

Implementation Details. We train the models using the original source code.³ For meta-training, we use a subset of a large-scale video detection dataset [38], and the 58 sequences from the VOT2013 [39], VOT2014 [40] and VOT2015 [41] datasets, excluding the sequences in the OTB2015 database, based on the same settings in the Meta-Tracker [19]. We apply the Gaussian DropGrad method

³ https://github.com/silverbottlep/meta_trackers.

with the dropout rate of 0.2. We use the default hyper-parameter settings and evaluate the performance with the models at the last training iteration.

Object Tracking Results. The results of online object tracking on the OTB2015 dataset are presented in Table 5. The one-pass evaluation (OPE) protocol without restarts at failures is used in the experiments. We measure the precision and success rate based on the center location error and the bounding-box overlap ratio, respectively. The precision is calculated with a threshold 20, and the success rate is the averaged value with the threshold ranging from 0 to 1 with a step of 0.05. We show that applying the proposed DropGrad method consistently improves the performance in precision and success rate on both MetaCREST and MetaSDNet trackers. We present sample results of object online tracking in Fig. 5. We apply the proposed DropGrad method on the MetaCREST and MetaSDNet methods and evaluate these models on the OTB2015 dataset. Compared with the original MetaCREST and MetaSDNet, models trained with the DropGrad method track objects more accurately.

Table 6. Viewpoint estimation results. The DropGrad method can be applied to few-shot viewpoint estimation frameworks to mitigate the overfitting problem.

Model	Acc30 (\uparrow)	MedErr (\downarrow)
MetaView [20]	45.00 \pm 0.45%	33.60 \pm 0.94 $^\circ$
MetaView w/ DropGrad	46.16 \pm 0.55%	33.10 \pm 0.82$^\circ$

4.4 Few-Shot Viewpoint Estimation

Viewpoint estimation aims to estimate the viewpoint (*i.e.*, 3D rotation), denoted as $R \in SO(3)$, between the camera and the object of a specific category in the image. Given a few examples (*i.e.*, 10 images in this work) of a novel category with viewpoint annotations, few-shot viewpoint estimation attempts to predict the viewpoint of arbitrary objects of the same category. In this problem, the support set D^s contains few images \mathbf{x}^s of a new class and the corresponding viewpoint annotations \mathbf{y}^s . We conduct experiments on the ObjectNet3D [42] dataset, a viewpoint estimation benchmark dataset which contains 100 categories. Using the same evaluation protocol in [20], we extract 76 and 17 categories for training and testing, respectively.

Implementation Details. We apply the regularization on the MetaView [20] method, which is a meta-Siamese viewpoint estimator that applies gradient-based adaptation for novel categories. We obtain the source code from the authors, and keep all the default setting for training. We apply the Gaussian DropGrad scheme with the dropout rate of 0.1. Since there is no validation set available, we pick the model trained in the last epoch for evaluation.

Viewpoint Estimation Results. We show the viewpoint estimation results in Table 6. The evaluation metrics include Acc30 and MedErr, which represent

the percentage of viewpoints with rotation error under 30° and the median rotation error, respectively. The overall performance is improved by applying the proposed DropGrad method to the MetaView model during training.

5 Conclusions

In this work, we propose a simple yet effective gradient dropout approach for regularizing the training of gradient-based meta-learning frameworks. The core idea is to impose uncertainty by augmenting the gradient in the adaptation step during meta-training. We propose two forms of noise regularization terms, including the Bernoulli and Gaussian distributions, and demonstrate that the proposed DropGrad improves the model performance in three learning tasks. In addition, extensive analysis and studies are provided to further understand the benefit of our method. One study on cross-domain few-shot classification is also conducted to show that the DropGrad method is able to mitigate the overfitting issue under a larger domain gap.

References

1. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017)
2. Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., Lillicrap, T.: Meta-learning with memory-augmented neural networks. In: ICML (2016)
3. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: NIPS (2017)
4. Gupta, A., Mendonca, R., Liu, Y., Abbeel, P., Levine, S.: Meta-reinforcement learning of structured exploration strategies. In: NeurIPS (2018)
5. Rakelly, K., Zhou, A., Quillen, D., Finn, C., Levine, S.: Efficient off-policy meta-reinforcement learning via probabilistic context variables. In: ICML (2019)
6. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: NeurIPS (2018)
7. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Meta-learning for domain generalization. In: AAAI (2018)
8. Kim, T., Yoon, J., Dia, O., Kim, S., Bengio, Y., Ahn, S.: Bayesian model-agnostic meta-learning. In: NeurIPS (2018)
9. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NIPS (2016)
10. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *JMLR* **15**, 1929–1958 (2014)
11. Wang, S., Manning, C.: Fast dropout training. In: ICML (2013)
12. Ghiasi, G., Lin, T.Y., Le, Q.V.: Dropblock: a regularization method for convolutional networks. In: NeurIPS (2018)
13. Larsson, G., Maire, M., Shakhnarovich, G.: Fractalnet: ultra-deep neural networks without residuals. In: ICLR (2017)
14. Tompson, J., Goroshin, R., Jain, A., LeCun, Y., Bregler, C.: Efficient object localization using convolutional networks. In: CVPR (2015)

15. Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., Fergus, R.: Regularization of neural networks using dropconnect. In: ICML (2013)
16. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: CVPR (2018)
17. Antoniou, A., Edwards, H., Storkey, A.: How to train your maml. In: ICLR (2019)
18. Li, Z., Zhou, F., Chen, F., Li, H.: Meta-sgd: Learning to learn quickly for few shot learning. arXiv preprint [arXiv:1707.09835](https://arxiv.org/abs/1707.09835) (2017)
19. Park, E., Berg, A.C.: Meta-tracker: fast and robust online adaptation for visual object trackers. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11207, pp. 587–604. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01219-9_35
20. Tseng, H.Y., et al.: Few-shot viewpoint estimation. In: BMVC (2019)
21. Welinder, P., et al.: Caltech-ucsd birds 200. Technical report CNS-TR-2010-001, California Institute of Technology (2010)
22. Rezende, D.J., Mohamed, S., Danihelka, I., Gregor, K., Wierstra, D.: One-shot generalization in deep generative models. JMLR **48**, (2016)
23. Oreshkin, B., López, P.R., Lacoste, A.: Tadam: task dependent adaptive metric for improved few-shot learning. In: NeurIPS (2018)
24. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: relation network for few-shot learning. In: CVPR (2018)
25. Finn, C., Xu, K., Levine, S.: Probabilistic model-agnostic meta-learning. In: NeurIPS (2018)
26. Rusu, A.A., et al.: Meta-learning with latent embedding optimization. In: ICLR (2019)
27. Ravi, S., Beaton, A.: Amortized bayesian meta-learning. In: ICLR (2019)
28. Goodfellow, I.J., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y.: Maxout networks. In: ICML (2013)
29. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: CVPR (2009)
30. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2017)
31. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: ICLR (2019)
32. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
33. Hilliard, N., Phillips, L., Howland, S., Yankov, A., Corley, C.D., Hodas, N.O.: Few-shot learning with metric-agnostic conditional embeddings. arXiv preprint [arXiv:1802.04376](https://arxiv.org/abs/1802.04376) (2018)
34. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: CVPR (2016)
35. Song, Y., Ma, C., Gong, L., Zhang, J., Lau, R.W.H., Yang, M.H.: Crest: convolutional residual learning for visual tracking. In: ICCV (2017)
36. Nam, H., Han, B.: Learning multi-domain convolutional neural networks for visual tracking. In: CVPR (2016)
37. Wu, Y., Lim, J., Yang, M.H.: Object tracking benchmark. TPAMI (2015)
38. Russakovsky, O., et al.: Imagenet large scale visual recognition challenge. IJCV (2015)
39. Kristan, M., et al.: T.V.: The visual object tracking vot2013 challenge results. In: ICCV Workshop (2013)
40. Kristan, M., et al.: A.L.: The visual object tracking vot2014 challenge results. In: ECCV Workshop (2014)

41. Kristan, M., et al., G.N.: The visual object tracking vot2015 challenge results. In: ICCV Workshop (2015)
42. Xiang, Yu., et al.: ObjectNet3D: a large scale database for 3D object recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV 2016. LNCS, vol. 9912, pp. 160–176. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_10