

Meta-Learning with Task-Adaptive Loss Function for Few-Shot Learning

Sungyong Baik¹ Janghoon Choi¹ Heewon Kim¹ Dohee Cho¹ Jaesik Min² Kyoung Mu Lee¹

¹ASRI, Department of ECE, Seoul National University ²Hyundai Motor Group

¹{dsybaik, ultio791, ghimhw, jdh12245, kyoungmu}@snu.ac.kr ²jaesik.min@hyundai.com

Abstract

In few-shot learning scenarios, the challenge is to generalize and perform well on new unseen examples when only very few labeled examples are available for each task. Model-agnostic meta-learning (MAML) has gained the popularity as one of the representative few-shot learning methods for its flexibility and applicability to diverse problems. However, MAML and its variants often resort to a simple loss function without any auxiliary loss function or regularization terms that can help achieve better generalization. The problem lies in that each application and task may require different auxiliary loss function, especially when tasks are diverse and distinct. Instead of attempting to hand-design an auxiliary loss function for each application and task, we introduce a new meta-learning framework with a loss function that adapts to each task. Our proposed framework, named Meta-Learning with Task-Adaptive Loss Function (MeTAL), demonstrates the effectiveness and the flexibility across various domains, such as few-shot classification and few-shot regression.

1. Introduction

Training deep neural networks entails a tremendous amount of labeled data and the corresponding efforts, which hinder the prompt application to new domains. As such, there has been growing interests in few-shot learning, in which the goal is to imbue the artificial intelligence systems with the capability of learning new concepts, given only few labeled examples (*e.g.* support examples). The core challenge in few-shot learning is to alleviate the susceptibility of deep neural networks to overfitting under few-data regime and achieve generalization on new examples (*e.g.* query examples).

Recently, meta-learning [34, 40], *a.k.a.* learning-to-learn, has emerged as one of the prominent methods for few-shot learning. Meta-learning is used in the field of few-shot learning to learn a learning framework that can adapt to novel tasks and generalize under few-data regime.

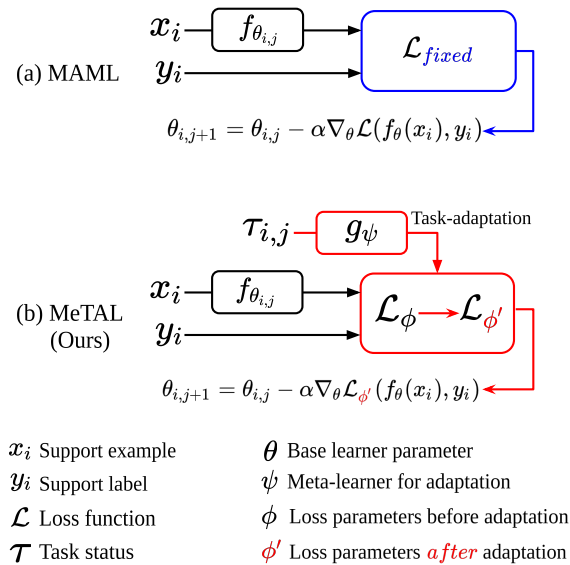


Figure 1. Overview of inner-loop optimization in optimization-based meta-learning frameworks. (a) Conventional approaches, such as MAML [10], utilize a fixed given classical loss function (*e.g.* cross entropy for classification) during adaptation to a task. (b) Our proposed scheme, MeTAL, instead meta-learns a loss function whose parameters ϕ are made adaptive to the current task state τ at the j -th step of adaptation to the i -th task.

Among the meta-learning algorithms, optimization-based meta-learning has enjoyed the attention from different domains for its flexibility that enables application across diverse domains. Optimization-based meta-learning algorithms are often formulated as bi-level optimization [10, 25, 29]. In such formulation, an outer-loop optimization trains a learning algorithm to achieve generalization while an inner-loop optimization uses the learning algorithm to adapt a base learner to a new task with few examples.

Model-agnostic meta-learning (MAML) [10], one of the seminal optimization-based meta-learning methods, learns an initial set of values of network weights to achieve generalization. The learned initialization serves as a good starting point for adapting to new tasks with few examples and few updates. Although the learned initialization is

trained to be a good starting point, MAML often faces the difficulty to achieve generalization, especially when tasks are diverse or significantly different between training and test phases [9]. Several works attempted to overcome the difficulty either by attempting to find a better initialization [5, 13, 11, 15, 43, 47] or a better fast adaptation process (inner-loop update rule) [2, 7, 19, 20, 32]. However, these methods resort to a simple loss function (*e.g.* cross-entropy in classification) in the inner-loop optimization even though other auxiliary loss functions, such as ℓ_2 regularization terms, can help achieve better generalization [4].

On the other hand, we focus on designing a better loss function for the inner-loop optimization in MAML framework. As outlined in Figure 1, we propose a new framework called **Meta-Learning with Task-Adaptive LLoss Function (MeTAL)** to learn an adaptive loss function that leads to better generalization for each task. Specifically, MeTAL learns a task-adaptive loss function via two meta-learners: one meta-learner for learning a loss function and one meta-learner for generating parameters that transform a learned loss function. Our task-adaptive loss function is designed to be flexible in that both labeled (*e.g.* support) and unlabeled (*e.g.* query) examples can be used together to adapt a base learner to each task during the inner-loop optimization.

The experimental results demonstrate that MeTAL greatly improves the generalization of MAML. Owing to the simplicity and flexibility of MeTAL, we further demonstrate its effectiveness not only across different domains but also other MAML-based algorithms. When applied to other MAML-based algorithms, MeTAL consistently brings a substantial boost in generalization performance, introducing a new state-of-the-art performance among MAML-based algorithms. This alludes to the significance of a task-adaptive loss function, which has drawn less attention in contrast to initialization schemes and inner-loop update rules. Overall experimental results underline that learning a better loss function for the inner loop optimization is important complementary component to learning a better inner-loop update or a better initialization.

2. Related Work

Few-shot learning aims to address scenarios where only few examples are available for each task. The ultimate goal is to learn new tasks with these given few examples while achieving generalization over unseen examples. To this end, meta-learning algorithms attempt to tackle the few-shot learning problem via the learning of prior knowledge from previous tasks, which is then used to adapt to new tasks without overfitting [8, 14, 34, 35, 40].

Depending on how the learning of prior knowledge and task adaptation process are formulated, meta-learning systems can generally be classified into metric-learning-based, black-box or network-based, and optimization-based ap-

proaches. Metric-learning-based approaches encode prior knowledge into an embedding space where similar (different) classes are closer (further apart) from each other [21, 17, 37, 38, 42]. Black-box or network-based approaches employ a network or external memory to directly generate weights [23, 24], weight updates [1, 14, 29], or predictions [22, 33]. Meanwhile, optimization-based approaches employ bi-level optimization to learn the learning procedures, such as initialization and weight updates, that will be used to adapt to new tasks with few examples [2, 4, 5, 10, 20, 25, 36].

In this work, we focus on model-agnostic meta-learning (MAML) algorithm [10], one of the most popular instances of the optimization approaches, owing to its simplicity and applicability across diverse problem domains. MAML formulates prior knowledge as a learnable initialization, from which good generalization performance can be achieved for a new task after gradient-based fine-tuning with the given few examples. While MAML is known for its simplicity and flexibility, it is also known for its relatively low generalization performance. There has been recent studies on improving the overall performance either by enhancing the learning scheme of initialization [5, 15, 32, 41, 43] or improving gradient-based fine-tuning process [2, 4, 20, 36].

However, the aforementioned works still employ only a common loss function corresponding to a task (*e.g.* cross-entropy in classification) during the inner-loop optimization. Common deep learning frameworks, on the other hand, often use auxiliary loss terms, such as ℓ_2 regularization term, in an effort to prevent overfitting. As the goal of few-shot learning is to achieve generalization over unseen examples after adaptation with only few examples, using auxiliary loss terms like a natural choice. Few recently introduced methods have applied auxiliary loss functions in the inner-loop optimization to reduce the computational cost [18, 28] or improve the generalization [12]. Other works attempted to learn loss functions for reinforcement learning (RL) [39, 45, 46, 48], supervised learning [6], and incorporating unsupervised learning into few-shot learning [3]. However, loss functions from these methods either have task-specific requirements, such as environment interaction in RL, or remain fixed after training. Fixed loss functions may be disadvantageous when new tasks may prefer different loss functions, especially in case train tasks and new tasks are significantly different (*e.g.* cross-domain few-shot classification [9]).

To this end, we propose a new meta-learning framework with a task-adaptive loss function (MeTAL). In particular, a task-specific loss function is learned by a meta-network whose parameters are adapted to the given task. Not only does MeTAL achieve outstanding performance, but it also maintains the simplicity and can be used jointly with other meta-learning algorithms.

Algorithm 1 Meta-learning with task-adaptive loss

Require: Task distribution $p(\mathcal{T})$
Require: Learning rates α, η
Require: Base learner network f , meta-networks g, l

- 1: Randomly initialize θ, ϕ, ψ
- 2: **while** not converged **do**
- 3: Sample a batch of tasks $\mathcal{T}_i \sim p(\mathcal{T})$
- 4: **for** each task \mathcal{T}_i **do**
- 5: Sample support set $\mathcal{D}_i^S = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{s=1}^K$ from \mathcal{T}_i
- 6: Sample query set $\mathcal{D}_i^Q = \{(\mathbf{x}_i^q, \mathbf{y}_i^q)\}_{q=1}^M$ from \mathcal{T}_i
- 7: Initialize $\theta_{i,0} = \theta$
- 8: **for** j in number of inner-loop updates J **do**
- 9: Adapt $\theta_{i,j+1} \leftarrow \theta_{i,j}$ using Algorithm 2
- 10: **end for**
- 11: Compute the loss on the query set:
 $\mathcal{L}(\mathcal{D}_i^Q; \theta_{i,j}) = \mathcal{L}(f(\mathbf{x}_i^q; \theta_{i,j}), \mathbf{y}_i^q)$
- 12: **end for**
- 13: Perform gradient descent to update weights:
 $(\theta, \phi, \psi) \leftarrow (\theta, \phi, \psi) - \eta \nabla_{(\theta, \phi, \psi)} \sum_{\mathcal{T}_i} \mathcal{L}(\mathcal{D}_i^Q; \theta_{i,j})$
- 14: **end while**

3. Proposed Method

3.1. Preliminaries

3.1.1 Problem formulation

We first introduce the preliminaries on meta-learning in the context of few-shot learning. The meta-learning framework assumes a collection of tasks $\{\mathcal{T}_i\}_{i=1}^T$, each of which is assumed to be drawn from a task distribution $p(\mathcal{T})$. Each task \mathcal{T}_i consists of two disjoint sets of a dataset \mathcal{D}_i : a support set \mathcal{D}_i^S and a query set \mathcal{D}_i^Q . Each set, in turn, consists of a number of pairs of input \mathbf{x} and output \mathbf{y} : $\mathcal{D}_i^S = \{(\mathbf{x}_i^s, \mathbf{y}_i^s)\}_{s=1}^K$ and $\mathcal{D}_i^Q = \{(\mathbf{x}_i^q, \mathbf{y}_i^q)\}_{q=1}^M$.

The goal of meta-learning is to learn a learning algorithm (formulated by a model with parameters ϕ) that can quickly learn tasks from a task distribution $p(\mathcal{T})$. The learned learning algorithm is then leveraged to learn a new task \mathcal{T}_i by adapting a base learner, parameterized by θ , using the task support examples \mathcal{D}_i^S , which is given by:

$$\theta_i = \arg \min_{\theta} \mathcal{L}(\mathcal{D}_i^S; \theta, \phi), \quad (1)$$

where \mathcal{L} denotes a loss function that evaluates the performance on a task. As the support set \mathcal{D}_i^S is used to learn a task, few-shot learning is often called k -shot learning when k number of support examples is available for each task ($|\mathcal{D}_i^S| = K = k$).

The resultant task-specific base learner is represented by parameters θ_i . The learned learning algorithm parameterized by ϕ is then evaluated on how a task-specific base learner θ_i generalizes to unseen query examples \mathcal{D}_i^Q . Thus,

Algorithm 2 Inner-loop update subroutine

Require: Weights $\theta_{i,j}$ of the base learner f
Require: Meta-networks g, l with parameters ϕ
Require: Support set examples \mathcal{D}_i^S
Require: Unlabeled query set examples $\{\mathbf{x}_i^q\}_{q=1}^M$ **if** *semi-supervised* setting

- 1: Compute the base learner output on the support set
 $f(\mathbf{x}_i^{s=1:K}; \theta_{i,j}) = \{f(\mathbf{x}_i^s; \theta_{i,j})\}_{s=1}^K$
- 2: Compute the loss on the support set:
 $\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j}) = \mathcal{L}(f(\mathbf{x}_i^{s=1:K}; \theta_{i,j}), \mathbf{y}_i^s)$
- 3: Compute the task state:
- 4: **if** *supervised* **then**
- 5: $\tau_{i,j} = [\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j}), \theta_{i,j}, \{f(\mathbf{x}_i^{s=1:K}; \theta_{i,j})\}_{s=1}^K]$
- 6: **else if** *semi-supervised* **then**
- 7: Compute the base learner output on the query set
 $f(\mathbf{x}_i^{q=1:M}; \theta_{i,j}) = \{f(\mathbf{x}_i^q; \theta_{i,j})\}_{q=1}^M$
- 8: $\tau_{i,j} = [\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j}), \theta_{i,j}, f(\mathbf{x}_i^{s=1:K}; \theta_{i,j}), f(\mathbf{x}_i^{q=1:M}; \theta_{i,j})]$
- 9: **end if**
- 10: Compute the affine transformation parameters:
 $\gamma_{i,j}, \beta_{i,j} = g(\tau_{i,j}; \psi)$
- 11: Adapt the loss function parameters:
 $\phi'_{i,j} = \gamma_{i,j} \phi + \beta_{i,j}$
- 12: Compute task-adaptive loss: $\mathcal{L}_{\phi'_{i,j}}(\tau_{i,j})$
- 13: Perform gradient descent to adapt f to \mathcal{T}_i :
 $\theta_{i,j+1} = \theta_{i,j} - \alpha \nabla_{\theta_{i,j}} \mathcal{L}_{\phi'_{i,j}}(\tau_{i,j})$

the objective of the meta-learning algorithm becomes,

$$\phi^* = \arg \min_{\phi} \mathbb{E}_{\mathcal{T}_i \sim p(\mathcal{T})} [\mathcal{L}(\mathcal{D}_i^Q; \theta_i, \phi)]. \quad (2)$$

3.1.2 Model-agnostic meta-learning

MAML [10] encodes prior knowledge into a learnable initialization that serves as a good initial set of values for weights of a base learner network across tasks. This formulation, in which meta-learning initialization for a base learner, leads to bi-level optimization: inner-loop optimization and outer-loop optimization. For the inner-loop optimization, a base learner is fine-tuned with support examples \mathcal{D}_i^S , from the learnable initialization θ , to each task for a fixed number of weight updates via gradient descent. Thus, after initializing $\theta_{i,0} = \theta$, the task adaptation objective (Equation (1)) is minimized via gradient-descent. The inner-loop optimization at j -th step is denoted as:

$$\theta_{i,j+1} = \theta_{i,j} - \alpha \nabla_{\theta_{i,j}} \mathcal{L}(\mathcal{D}_i^S; \theta_{i,j}). \quad (3)$$

Then, after J number of inner-loop update steps, the task-specific base learner parameters θ_i become $\theta_{i,J}$.

In the case of the outer-loop optimization, the meta-learned initialization θ is evaluated by the generalization

performance of the task-specific base learner with parameters θ_i (or $\theta_{i,j}$) on unseen query examples \mathcal{D}_i^Q . The evaluated generalization on unseen examples is then used as a feedback signal to update the initialization θ . In other words, MAML minimizes the objective of the meta-learning algorithm, as in Equation (2), as follows:

$$\theta \leftarrow \theta - \eta \nabla_{\theta} \sum_{\mathcal{T}_i} \mathcal{L}(\mathcal{D}_i^Q; \theta_i). \quad (4)$$

3.2. Meta-learning with Task-Adaptive Loss Function (MeTAL)

3.2.1 Overview

Previous meta-learning formulations assume a fully-supervised setting for a given task \mathcal{T}_i where they use labeled examples in the support set \mathcal{D}_i^S to find the task-specific base learner θ_i through minimizing a fixed given loss function \mathcal{L} . On the other hand, we aim to control or meta-learn a loss function itself that would regulate the whole adaptation or inner-loop optimization process for better generalization. We start with meta-learning an inner-loop optimization loss function $\mathcal{L}_{\phi}(\cdot)$, modeled by a small neural network with its meta-learnable parameters ϕ . Thus, the inner-loop update in Equation (3) becomes,

$$\theta_{i,j+1} = \theta_{i,j} - \alpha \nabla_{\theta_{i,j}} \mathcal{L}_{\phi}(\tau_{i,j}), \quad (5)$$

where $\tau_{i,j}$ denotes the task state for \mathcal{T}_i at time-step j , which is usually just the support set \mathcal{D}_i^S in the case of the typical meta-learning formulation, as in Equation (3). As different tasks (especially under cross-domain scenarios [9]) may prefer different regularization or auxiliary loss functions or even loss functions itself during adaptation to achieve better generalization, we aim to learn to adapt a loss function itself to each task. To enable a meta-learned loss function to be adaptive, one natural design choice could be to perform gradient descent, similar to how base learner parameters θ_i are updated as in Equation (3). However, such design would result in a large computation graph, especially if a meta-learning algorithm is trained with higher-order gradients. Alternatively, affine transformation could be applied to make a loss function adaptive to the given task. Affine transformation conditioned on some input has been proved to be effective by several works in making feature responses adaptive [27, 26, 16] and making meta-learned initialization adaptive [43]. In order to make a loss function task-adaptive without huge computation burden, we propose to dynamically transform loss function parameters ϕ via affine transformation:

$$\phi' = \gamma \phi + \beta, \quad (6)$$

where ϕ are the meta-learnable loss function parameters and γ, β are the transformation parameters generated by the meta-learner $g(\tau_j; \psi)$ which is parameterized by ψ .

To train our meta-learning framework to generalize across different tasks, which involves optimizing the parameters θ, ϕ , and ψ , the outer-loop optimization is performed with each task \mathcal{T}_i given the respective task-specific learner θ_i and its examples in the query set \mathcal{D}_i^Q as in,

$$(\theta, \phi, \psi) \leftarrow (\theta, \phi, \psi) - \eta \nabla_{(\theta, \phi, \psi)} \sum_{\mathcal{T}_i} \mathcal{L}(\mathcal{D}_i^Q; \theta_i). \quad (7)$$

The overall training procedure of our method is summarized in Algorithm 1.

3.2.2 Task-adaptive loss function

Since our loss meta-learner \mathcal{L}_{ϕ} and the meta-learner g_{ψ} are modeled using neural networks, their inputs can be formulated to contain auxiliary task-specific information on the intermediate learning status, which we define as a task state τ . At j -th inner-loop step for a given task \mathcal{T}_i , in addition to classical loss information $\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j})$ (evaluated on the labeled support set examples \mathcal{D}_i^S), auxiliary learning state information, such as the network weights $\theta_{i,j}$ and the output values $f(\mathbf{x}_i^s; \theta_{i,j})$, can be included in the task state $\tau_{i,j}$.

Moreover, we can also include the base learner responses on the unlabeled examples \mathbf{x}_i^q from the query set in the task state, which enables the inner-loop optimization to perform *semi-supervised* learning. This shows that our framework can use such additional task-specific information for fast adaptation, which was rarely utilized in previous MAML-based meta-learning algorithms, whereas metric-based meta-learning algorithms, such as [21], attempt to utilize unlabeled query examples to maximize the performance. The semi-supervised inner-loop optimization maximizes the advantage of transductive setting (all query examples are assumed to be available at once), which MAML-based algorithms have implicitly used for better performance [25]. The procedure of inner-loop optimization with the task-adaptive loss function for both supervised and semi-supervised setting is organized in Algorithm 2.

3.2.3 Architecture

For our task-adaptive loss function \mathcal{L}_{ϕ} , we employ a 2-layer MLP with ReLU activation between the layers, which returns a single scalar value as output. For the improved computational efficiency, the task state $\tau_{i,j}$ used in the inner-loop optimization is formulated as a concatenation of mean of support set losses $\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j})$, layer-wise means of base learner weights $\theta_{i,j}$, and example-wise mean of base learner output values $f(\mathbf{x}_i^s; \theta_{i,j})$. Assuming a L -layer neural network for the base learner f which returns N -dimensional output values (for N -way classification), the dimension of the task state $\tau_{i,j}$ becomes $1 + L + N$, which is computationally minimal. This can slightly increase under the semi-supervised learning setting where additional information can be derived from the responses of a base learner $f(\mathbf{x}_i^q; \theta_{i,j})$ on the unlabeled query examples.

Meta-network g_ψ also employs a 2-layer MLP with ReLU activation between the layers. The network produces layer-wise affine transformation parameters that are applied to the loss function parameters ϕ . Since our meta-learning framework does not impose any constraint on the base learner f and its target applications, our formulation is general and can be readily applied to any gradient-based differentiable learning algorithm. For more details, please refer to the supplementary document and our code¹.

4. Experiments

In this section, we perform experiments on several few-shot learning problems, such as few-shot classification, cross-domain few-shot classification, and few-shot regression to corroborate the effectiveness of task-adaptive loss functions. All experimental results by our proposed method MeTAL are performed with semi-supervised inner-loop optimization, in which labeled support examples and unlabeled query examples are used together for the inner-loop optimization. Note that we do not use extra data and that MeTAL simply takes more benefits from a transductive setting (all query examples are available at once), which other MAML-variants also employ for higher performance [25].

4.1. Few-shot classification

In few-shot classification, each task is defined as N -way k -shot classification, in which N is the number of classes and k is the number of examples (shots) per each class.

4.1.1 Datasets

The most commonly used datasets for few-shot classification are two ImageNet[31]-derivative datasets: miniImageNet [29] and tieredImageNet [30]. Both datasets are composed of three disjoint subsets (train, validation, and test sets), each of which consists of images with the size of 84×84 . The datasets differ in how classes are split into disjoint subsets. miniImageNet randomly samples and groups classes into 64 classes for meta-training, 16 for meta-validation, and 20 for meta-test [29]. tieredImageNet, on the other hand, groups classes into 34 categories according to ImageNet class hierarchy and splits groups into 20 categories for meta-training, 6 for meta-validation, and 8 for meta-test [30] in an effort to minimize class similarities between the three disjoint sets.

4.1.2 Results

We assess our method MeTAL and compare with other MAML variants on miniImageNet and tieredImageNet under two typical settings: 5-way 5-shot and 5-way 1-shot

¹The code is available at <https://github.com/baiksung/MeTAL>

classification, as presented in Table 1. The results demonstrate that not only does MeTAL greatly improve the generalization performance of MAML but also can be applied in conjunction with other MAML variants, such as MAML++ [2] and ALFA [4], to bring further improvement. MAML++ learns fixed step-and-layer-wise inner-loop learning rates while ALFA learns task-adaptive inner-loop learning rates and regularization terms. Although these methods do not consider a loss function to be learnable, if a loss function is regarded as a part of the model, then MeTAL may be seen as a more general extension of these methods. However, the further improvement brought by MeTAL upon these methods demonstrates that improving the inner-loop optimization objective function is not just a simple extension but a complementary and orthogonal factor. The major contribution of MeTAL lies in formulating an inner-loop loss function to be learnable and task-adaptive. Furthermore, MeTAL, together with ALFA [4], greatly outperforms other models that either use larger networks, such as DenseNet or WideResNet, or are pretrained or trained with data augmentation. These results suggest the effectiveness of our learned task-adaptive loss function in achieving better generalization.

4.2. Cross-domain few-shot classification

Cross-domain few-shot classification, introduced by Chen *et al.* [9], tackles a more challenging and practical few-shot classification scenario, where meta-train tasks and meta-test tasks are sampled from different task distributions. Such scenario is intentionally designed to create a large domain gap between meta-train and meta-test, thereby assessing the susceptibility of meta-learning algorithms to meta-level overfitting. Specifically, a meta-learning algorithm can be said to be meta-overfitted if the algorithm is too dependent on prior knowledge from previously seen meta-train tasks, instead of focusing on the given few examples to learn a new task. This meta-level overfitting will result in a learning system being more likely to fail to adapt to a new task that is sampled from substantially different task distributions.

4.2.1 Datasets

To simulate such challenging scenario, Chen *et al.* [9] first meta-trains algorithms on miniImageNet [29] and evaluates them on CUB dataset (CUB-200-2011) [44] during meta-test. In contrast to ImageNet that is compiled for general classification tasks, CUB targets fine-grained classification. Following the protocol from [9], 200 classes of the dataset are split into 100 meta-train, 50 meta-validation, and 50 meta-test sets.

4.2.2 Results

Table 2 presents the performance of MAML [10], one of recent MAML variants ALFA [4], and MeTAL when they

Model	Base learner Backbone	miniImageNet		tieredImageNet	
		1-shot	5-shot	1-shot	5-shot
MAML + L2F [5]	4-CONV	52.10 ± 0.50%	69.38 ± 0.46%	54.40 ± 0.50%	73.34 ± 0.44%
MAML [10]	4-CONV	48.70 ± 1.75%	63.11 ± 0.91%	-	-
MAML [‡]	4-CONV	49.64 ± 0.31%	64.99 ± 0.27%	50.98 ± 0.26%	66.25 ± 0.19%
MeTAL (Ours)	4-CONV	52.63 ± 0.37%	70.52 ± 0.29%	54.34 ± 0.31%	70.40 ± 0.21%
MAML++ + SCA [3]	4-CONV	54.24 ± 0.99%	71.85 ± 0.53%	-	-
MAML++ [2]	4-CONV	52.15 ± 0.26%	68.32 ± 0.44%	-	-
MAML++ + MeTAL (Ours)	4-CONV	57.18 ± 0.42%	72.89 ± 0.44%	59.93 ± 0.36%	75.39 ± 0.29%
ALFA + MAML [4]	4-CONV	50.58 ± 0.51%	69.12 ± 0.47%	53.16 ± 0.49%	70.54 ± 0.46%
ALFA + MeTAL (Ours)	4-CONV	57.75 ± 0.38%	74.10 ± 0.43%	60.29 ± 0.37%	75.88 ± 0.29%
MAML++ [2]	DenseNet	58.37 ± 0.27%	75.50 ± 0.19%	-	-
SCA + MAML++ [3]	DenseNet	62.86 ± 0.79%	77.64 ± 0.40%	-	-
MAML + L2F [5]	ResNet12	59.71 ± 0.49%	77.04 ± 0.42%	64.04 ± 0.48%	81.13 ± 0.39%
MAML [‡]	ResNet12	58.60 ± 0.42%	69.54 ± 0.38%	59.82 ± 0.41%	73.17 ± 0.32%
MeTAL (Ours)	ResNet12	59.64 ± 0.38%	76.20 ± 0.19%	63.89 ± 0.43%	80.14 ± 0.40%
ALFA + MAML [4]	ResNet12	59.74 ± 0.49%	77.96 ± 0.41%	64.62 ± 0.49%	82.48 ± 0.38%
ALFA + MeTAL (Ours)	ResNet12	66.61 ± 0.28%	81.43 ± 0.25%	70.29 ± 0.40%	86.17 ± 0.35%
SNAIL [22]	ResNet12	55.71 ± 0.99%	68.88 ± 0.92%	-	-
AdaResNet [24]	ResNet12	56.88 ± 0.62%	71.94 ± 0.57%	-	-
TADAM [26]	ResNet12	58.50 ± 0.30%	76.70 ± 0.30%	-	-
LEO-trainval* [32]	WRN-28-10	61.76 ± 0.08%	77.59 ± 0.12%	66.33 ± 0.05%	81.44 ± 0.09%
MetaOpt [†] [18]	ResNet12	62.64 ± 0.61%	78.63 ± 0.46%	65.99 ± 0.72%	81.56 ± 0.53%

* Pretrained

† Trained with data augmentation.

‡ Reproduced.

Table 1. 5-way 1-shot and 5-way 5-shot classification test accuracy on miniImageNet and tieredImageNet.

	Base learner Backbone	miniImageNet → CUB
MAML	4-CONV	52.70 ± 0.32%
MeTAL (Ours)	4-CONV	58.20 ± 0.24%
ALFA + MAML	4-CONV	58.35 ± 0.25%
ALFA + MeTAL (Ours)	4-CONV	66.37 ± 0.17%
MAML	ResNet12	53.83 ± 0.32%
MeTAL (Ours)	ResNet12	61.29 ± 0.21%
ALFA + MAML	ResNet12	63.64 ± 0.42%
ALFA + MeTAL (Ours)	ResNet12	70.22 ± 0.14%

Table 2. 5-way 5-shot cross-domain few-shot classification. Models are trained on miniImageNet and evaluated on CUB.

are trained on miniImageNet meta-train set and evaluated on CUB meta-test set. Similar to the few-shot classification results outlined in Table 1, MeTAL is shown to greatly improve the generalization even under the more challenging cross-domain few-shot classification scenario. In fact, MeTAL improves the performance of both MAML and ALFA + MAML by greater extent in cross-domain few-shot classification ($\sim 8\%$) than in few-shot classification

($\sim 4\%$). This implies the effectiveness of MeTAL in learning new tasks from different domains and its robustness to the domain gap, stressing the importance of task-adaptive loss function. Another observation can be made regarding the results: the increase in generalization performance made by MeTAL on ALFA + MAML is as great as on MAML, indicating the orthogonality of the problem MeTAL attempts to solve. ALFA [4] also aims to improve the inner-loop optimization but the difference is they focus on developing a new weight-update rule (gradient descent). On the other hand, we focus on a loss function that is used in the inner-loop optimization. The consistent generalization improvement by MeTAL across different baselines and architectures signifies that designing a better inner-loop optimization loss function is important factor and complementary to designing a better weight-update rule.

4.3. Few-shot regression

To demonstrate the flexibility and applicability of our method MeTAL, we evaluate MAML and MeTAL on few-shot regression, or k -shot regression. In k -shot regression,

	5 shots	10 shots	20 shots
MAML	0.86 ± 0.23	0.50 ± 0.12	0.26 ± 0.08
MeTAL (Ours)	0.74 ± 0.18	0.44 ± 0.11	0.21 ± 0.06

Table 3. k -shot regression: mean-square error (MSE) measured over 100 sampled points with 95% confidence intervals.

	cross entropy	learned loss	accuracy
(1)	✓		$64.99 \pm 0.27\%$
(2)		✓	$67.42 \pm 0.34\%$
(3)	✓	✓	$67.86 \pm 0.42\%$

Table 4. The study on the effectiveness of learning a loss function during the inner-loop optimization. Model (1) denotes MAML.

each task is to estimate a given unknown function when only very few number (k) of whose sampled points are given. The task distribution consists of tasks that have a target function with parameters whose values vary within a defined range. In this work, we follow the general settings Finn *et al.* [10] have used for evaluating MAML. Specifically, each task has a sinusoid $y(x) = A \sin(\omega x + b)$ as a target function whose parameter values are within the following ranges: amplitude $A \in [0.1, 5.0]$, frequency $\omega \in [0.8, 1.2]$, and phase $b \in [0, \pi]$. For each task, input data points x are sampled from $[-5.0, 5.0]$. Regression is performed by performing a single gradient descent on a base learner whose neural architecture is composed of 3 layers of size 80 with ReLU non-linear activation functions in between. The performance is measured in mean-square error (MSE) between the estimated output values \hat{y} and ground-truth output values y .

Table 3 outlines the regression results from MAML [10] and MeTAL under 5-shot, 10-shot, and 20-shot settings. Again, MeTAL demonstrates the consistent performance improvement across different settings. This proves the applicability and flexibility of the proposed task-adaptive loss function, learned by MeTAL.

4.4. Ablation studies

To investigate the contribution of each module in MeTAL, we perform ablation study experiments in this section. In particular, we analyze the effectiveness of the task state information, the learning of a loss function, the task-adaptive loss function, and the semi-supervised inner-loop optimization formulation. All ablation study experiments were performed with a base learner that has 4-CONV backbone under 5-way 5-shot few-shot classification.

4.4.1 Learning a loss function

First, we analyze the importance of learning an inner-loop optimization loss function. In detail, the performance is measured when the inner-loop optimization is performed with a loss function that is learned without the adaptation (*i.e.* only a meta-network \mathcal{L}_ϕ is used for model (2), (3)) and

	Task-adaptive	semi-supervised	accuracy
(2)			$67.42 \pm 0.34\%$
(4)	✓		$68.56 \pm 0.36\%$
(5)		✓	$68.24 \pm 0.31\%$
(6)	✓	✓	$70.52 \pm 0.29\%$

Table 5. The study on the effectiveness of task-adaptive loss function and semi-supervised setting during the inner-loop optimization. Model (6) corresponds to MeTAL. Model (2) refers to MAML trained with a fixed learned loss from Table 4.

	cross entropy (support set)	weight	prediction	accuracy
(7)	✓			$67.86 \pm 0.41\%$
(8)	✓	✓		$68.66 \pm 0.46\%$
(9)	✓		✓	$67.94 \pm 0.47\%$
(6)	✓	✓	✓	$70.52 \pm 0.29\%$

Table 6. Investigation on the role of each factor of the task state τ .

compared with when a simple cross entropy is used (*i.e.* MAML is represented as model (1)).

The ablation study results summarized in Table 4 show that the learned a loss function helps MAML achieve better generalization, suggesting that a meta-learner has managed to learn loss function that is useful for generalization. Further, when cross entropy and a learned loss are used together, there is no significant difference from when only a learned loss is used, implying that a learned loss is able to maintain cross entropy loss information that is fed as input.

4.4.2 Task-adaptive loss function

We then examine the influence of task-adaptive loss function on the overall proposed framework. To this end, we use a meta-model g_ψ to generate affine transformation parameters, which are then used to adapt the parameters of a loss function meta-network \mathcal{L}_ϕ in model (2) from Table 4, according to Equation (6). The derived meta-learning algorithm, which is MeTAL without semi-supervised inner-loop optimization, is denoted as model (4) in Table 5. As shown in the table, the meta-learning algorithm benefits from a task-adaptive learned loss function, compared to a fixed learned function.

4.4.3 Semi-supervised inner-loop optimization

Next, we look into the effectiveness of the semi-supervised inner-loop optimization formulation. Similar to the task-adaptive loss function ablation study, we first derive a new model that is created by adding the semi-supervised inner-loop optimization formulation (using labeled support examples and unlabeled query examples together for fast adaptation via learned loss function) to model (2) from Table 4. Consequently, the resulting model, which is denoted as model (5), lacks the task-adaptive property, compared to our




Figure 2. The illustration of the generated affine transformation parameters, γ and β , that are generated by one of our proposed meta-networks, g . These values are then used to adapt the loss meta-network l to the given task. In particular, the figure visualizes the generated values for the weight of the first layer of the loss meta-network l . The generated values demonstrate its dynamic range across the inner-loop steps, which indicates a different loss function is preferred at each step. Additionally, different values are observed among different tasks, especially at the last inner-loop step, alluding to diverse preferences of tasks on the inner-loop optimization loss functions.

final method MeTAL. While the semi-supervised inner-loop optimization contributes to the performance improvement, it still lags behind the full algorithm MeTAL (denoted as model (6)), alluding to the significance of the task-adaptive loss function.

4.4.4 Task state

We perform another ablation study for the investigation on the effect of each factor of the task state τ : namely, the current weight values of base learner $\theta_{i,j}$, the output of network ($f(x_i^s; \theta_{i,j})$ for support and $f(x_i^q; \theta_{i,j})$ for query), and cross entropy loss for support set $\mathcal{L}(\mathcal{D}_i^S; \theta_{i,j}) = \mathcal{L}(f(x_i^s; \theta_i), \mathbf{y}_i^s)$. The ablation results are summarized in Table 6. When the original cross entropy on support examples is not included in the task state, the whole inner-loop optimization becomes unsupervised learning setting as no ground-truth information is involved during the inner-loop optimization. In such case, as one would expect, MeTAL struggles to achieve generalization, thereby excluding these results in the table. When conditioned on cross entropy loss as a task state (model (7)), MeTAL manages to bring better generalization. Furthermore, including weight (model (8)) or prediction (model (9)) into task state contributes to further improvement. Finally, when all factors of the task state are used, MeTAL achieves the best performance, emphasizing the importance of each factor.

4.5. Visualization

Figure 2 illustrates the affine transformation parameters γ and β generated by one of our proposed meta-networks g across tasks (represented as boxplot) for each inner-loop step. Observing how the generated γ and β values vary across the inner-loop steps, we can claim that MeTAL manages to dynamically adapt a loss function as the learning

state changes during the inner-loop optimization. Further, the generated parameter values are shown to vary among tasks, especially at the last inner-loop step. This may suggest that the overall framework is trained to make the most difference among tasks at the last step. Regardless, the dynamic range of the generated affine transformation parameter values among tasks validates the effectiveness of the MeTAL in adapting a loss function to the given task.

5. Conclusion

In this work, we propose a meta-learning framework with a task-adaptive loss function for few-shot learning. The proposed scheme, named MeTAL, learns a loss function that adapts to each task based on the current task state during the inner-loop optimization. Consequently, MeTAL is able to learn a loss function that is specifically needed by each task for better generalization. Furthermore, not only does the flexibility of MeTAL enable it to be applied to different MAML variants and problem domains, but also allows for the semi-supervised inner-loop optimization, in which labeled support examples and unlabeled query examples are used jointly for adapting to tasks. Overall, the experimental results underline the importance of learning a good loss function for each task, which has drawn relatively less attention, compared to a weight-update rule or initialization in the context of few-shot learning.

Acknowledgment This work was supported in part by IITP grant funded by the Korea government [No. 2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)], and in part by AIRS Company in Hyundai Motor and Kia through HMC/KIA-SNU AI Consortium Fund.

References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gómez, Matthew W. Hoffman, David Pfau, Tom Schaul, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In *NIPS*, 2016. 2
- [2] Antreas Antoniou, Harrison Edwards, and Amos Storkey. How to train your maml. In *ICLR*, 2019. 2, 5, 6
- [3] Antreas Antoniou and Amos Storkey. Learning to learn via self-critique. In *NeurIPS*, 2019. 2, 6
- [4] Sungyong Baik, Myungsub Choi, Janghoon Choi, Heewon Kim, and Kyoung Mu Lee. Meta-learning with adaptive hyperparameters. In *NeurIPS*, 2020. 2, 5, 6
- [5] Sungyong Baik, Seokil Hong, and Kyoung Mu Lee. Learning to forget for meta-learning. In *CVPR*, 2020. 2, 6
- [6] Sarah Bechtle, Artem Molchanov, Yevgen Chebotar, Edward Grefenstette, Ludovic Righetti, Gaurav Sukhatme, and Franziska Meier. Meta-learning via learned loss. In *ICPR*, 2021. 2
- [7] Harkirat Singh Behl, Atilim Günes Baydin, and Philip H.S. Torr. Alpha maml: Adaptive model-agnostic meta-learning. In *ICMLW*, 2019. 2
- [8] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, pages 6–8. Univ. of Texas, 1992. 2
- [9] Wei-Yu Chen, Yen-Cheng Liu, Zsolt Kira, Yu-Chiang Wang, and Jia-Bin Huang. A closer look at few-shot classification. In *ICLR*, 2019. 2, 4, 5
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017. 1, 2, 3, 5, 6, 7
- [11] Chelsea Finn, Kelvin Xu, and Sergey Levine. Probabilistic model-agnostic meta-learning. In *NeurIPS*, 2018. 2
- [12] Micah Goldblum, Steven Reich, Liam Fowl, Renkun Ni, Valeriia Cherepanova, and Tom Goldstein. Unraveling meta-learning: Understanding feature representations for few-shot tasks. In *ICML*, 2020. 2
- [13] Erin Grant, Chelsea Finn, Sergey Levine, Trevor Darrell, and Thomas Griffiths. Recasting gradient-based meta-learning as hierarchical bayes. In *ICLR*, 2018. 2
- [14] Sepp Hochreiter, A Younger, and Peter Conwell. Learning to learn using gradient descent. In *ICANN*, 2001. 2
- [15] Muhammad Abdullah Jamal and Guo-Jun Qi. Task agnostic meta-learning for few-shot learning. In *CVPR*, 2019. 2
- [16] Xiang Jiang, Mohammad Havaei, Farshid Varno, Gabriel Chartrand, Nicolas Chapados, and Stan Matwin. Learning to learn with conditional class dependencies. In *ICLR*, 2019. 4
- [17] Gregory Koch, Richard Zemel, and Ruslan Salakhutdinov. Siamese neural networks for one-shot image recognition. In *ICMLW*, 2015. 2
- [18] Kwonjoon Lee, Subhansu Maji, Avinash Ravichandran, and Stefano Soatto. Meta-learning with differentiable convex optimization. In *CVPR*, 2019. 2, 6
- [19] Yoonho Lee and Seungjin Choi. Gradient-based meta-learning with learned layerwise metric and subspace. In *ICML*, 2018. 2
- [20] Zhenguo Li, Fengwei Zhou, Fei Chen, and Hang Li. Meta-sgd: Learning to learn quickly for few shot learning. *arXiv preprint arXiv:1707.09835*, 2017. 2
- [21] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. In *ICLR*, 2019. 2, 4
- [22] Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. A simple neural attentive meta-learner. In *ICLR*, 2018. 2, 6
- [23] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *ICML*, 2017. 2
- [24] Tsendsuren Munkhdalai, Xingdi Yuan, Soroush Mehri, and Adam Trischler. Rapid adaptation with conditionally shifted neurons. In *ICML*, 2018. 2, 6
- [25] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 1, 2, 4, 5
- [26] Boris N. Oreshkin, Pau Rodriguez, and Alexandre Lacoste. Tadam: Task dependent adaptive metric for improved few-shot learning. In *NeurIPS*, 2018. 4, 6
- [27] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *AAAI*, 2018. 4
- [28] Aravind Rajeswaran, Chelsea Finn, Sham Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *NeurIPS*, 2019. 2
- [29] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *ICLR*, 2017. 1, 2, 5
- [30] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B. Tenenbaum, Hugo Larochelle, and Richard S. Zemel. Meta-learning for semi-supervised few-shot classification. In *ICLR*, 2018. 5
- [31] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 5
- [32] Andrei A. Rusu, Dushyant Rao, Jakub Sygnowski, Oriol Vinyals, Razvan Pascanu, Simon Osindero, and Raia Hadsell. Meta-learning with latent embedding optimization. In *ICLR*, 2019. 2, 6
- [33] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *ICLR*, 2016. 2
- [34] Jurgen Schmidhuber. Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook. Diploma thesis, Technische Universitat Munchen, Germany, 1987. 1, 2
- [35] Jürgen Schmidhuber. Learning to control fast-weight memories: An alternative to dynamic recurrent networks. *Neural Computation*, 1992. 2
- [36] Christian Simon, Piotr Koniusz, Richard Nock, and Mehrtash Harandi. On modulating the gradient for meta-learning. In *ECCV*, 2020. 2

- [37] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NIPS*, 2017. 2
- [38] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H.S. Torr, and Timothy M. Hospedales. Learning to compare: Relation network for few-shot learning. In *CVPR*, 2018. 2
- [39] Flood Sung, Li Zhang, Tao Xiang, Timothy Hospedales, and Yongxin Yang. Learning to learn: Meta-critic networks for sample efficient learning. *arXiv preprint arXiv:1706.09529*, 2017. 2
- [40] Sebastian Thrun and Lorien Pratt. *Learning to learn*. Springer Science & Business Media, 2012. 1, 2
- [41] Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, and Hugo Larochelle. Meta-dataset: A dataset of datasets for learning to learn from few examples. In *ICLR*, 2020. 2
- [42] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. In *NIPS*, 2016. 2
- [43] Risto Vuorio, Shao-Hua Sun, Hexiang Hu, and Joseph J. Lim. Multimodal model-agnostic meta-learning via task-aware modulation. In *NeurIPS*, 2019. 2, 4
- [44] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical report, Caltech, 2011. 5
- [45] Zhongwen Xu, Hado van Hasselt, Matteo Hessel, Junhyuk Oh, Satinder Singh, and David Silver. Meta-gradient reinforcement learning with an objective discovered online. In *NeurIPS*, 2020. 2
- [46] Zhongwen Xu, Hado van Hasselt, and David Silver. Meta-gradient reinforcement learning. In *NeurIPS*, 2018. 2
- [47] Huaxiu Yao, Ying Wei, Junzhou Huang, and Zhenhui Li. Hierarchically structured meta-learning. In *ICML*, 2019. 2
- [48] Yang Yu. Towards sample efficient reinforcement learning. In *IJCAI*, 2018. 2