
Gradient Agreement as an Optimization Objective for Meta-Learning

Amir Erfan Eshratifar,¹ David Eigen,² Massoud Pedram,¹

¹Department of Electrical Engineering at University of Southern California, Los Angeles, CA 90089, USA

²Clarifai, San Francisco, CA 94105, USA

eshratif@usc.edu, deigen@clarifai.com, pedram@usc.edu

Abstract

This paper presents a novel optimization method for maximizing generalization over tasks in meta-learning. The goal of meta-learning is to learn a model for an agent adapting rapidly when presented with previously unseen tasks. Tasks are sampled from a specific distribution which is assumed to be similar for both seen and unseen tasks. We focus on a family of meta-learning methods learning initial parameters of a base model which can be fine-tuned quickly on a new task, by few gradient steps (MAML). Our approach is based on pushing the parameters of the model to a direction in which tasks have more agreement upon. If the gradients of a task agree with the parameters update vector, then their inner product will be a large positive value. As a result, given a batch of tasks to be optimized for, we associate a positive (negative) weight to the loss function of a task, if the inner product between its gradients and the average of the gradients of all tasks in the batch is a positive (negative) value. Therefore, the degree of the contribution of a task to the parameter updates is controlled by introducing a set of weights on the loss function of the tasks. Our method can be easily integrated with the current meta-learning algorithms for neural networks. Our experiments demonstrate that it yields models with better generalization compared to MAML and Reptile.

1 Introduction

Many successes of deep learning today rely on enormous amounts of labeled data, which is not practical for problems with small data, in which acquiring enough labeled data is expensive and laborious. Moreover, in many mission-critical applications such as autonomous vehicles, and drones, an agent needs to adapt rapidly to unseen environments. The ability to learn rapidly is a key characteristic of humans that distinguishes them from artificial agents. Humans leverage prior knowledge learned earlier when adapting to a changing task, an ability that can be explained as Bayesian inference (Brenden M Lake and Tenenbaum [2015]). Pioneered by (Schmidhuber [1987]), meta-learning has emerged recently as a novel field of study for learning from small amounts of data. Meta-learning algorithms learn their base models by sampling many different smaller tasks from a large data source instead of trying to emulate a computationally intractable Bayesian inference. As a result, one might expect that the meta-learned model is capable of generalizing well to new unseen tasks because of the task-agnostic way of training.

In this paper, we focus on a family of meta-learning algorithms that aim to learn the initialization of a network, which is then fine-tuned at test time on the new task by few steps of gradient descent. The most prominent work in this family is MAML (Finn et al. [2017]). MAML directly optimizes the performance of the model with respect to the test time fine-tuning procedure leading to a better initialization point.

In the parallel or batched version of MAML style algorithms (e.g. Reptile (Nichol et al. [2018]), at each iteration we optimize on N tasks. However, the contribution of each task to the model parameter updates is assumed to be the same, i.e., the update rule $\theta \leftarrow \theta - \beta \nabla \sum L_{\tau_i}(f_{\theta_i})$ is merely the average of the tasks losses. As an example, suppose a scenario in which a batch of N tasks is sampled and

2nd Workshop on Meta-Learning at NeurIPS 2018, Montréal, Canada.

$N - 1$ tasks have strong agreement on the gradient direction but one task has a large disagreeing gradient. In this case, the disagreeing task can adversely affect the next model parameter update.

The contribution of this paper is introducing an optimization objective, i.e., gradient agreement by which a model can generalize better across tasks. The proposed algorithm and its mathematical analysis are illustrated in Sections 2 and 3, respectively. In Section 4, we provide some empirical insights by evaluations on miniImageNet (Vinyals et al. [2016]) and Omniglot (Lake et al. [2011]) classification tasks, and a toy example regression task.

2 Gradient Agreement

In this work, we consider the optimization problem of MAML which is finding a good initial set of parameters, θ . The goal is that for a randomly sampled task τ with corresponding loss L_τ , the model converges to a low loss point after k number of gradient steps. This objective can be presented as $\min_{\theta} \mathbb{E}[L_\tau(U_\tau^k(\theta))]$, where $U_\tau^k(\theta)$ is a function performing k gradient steps that updates θ by the samples drawn from task τ . Minimizing this objective function requires two steps in MAML style works: 1) Adapting to a new task T_i in which the model parameters θ become θ_i by k steps of gradient descent on task T_i , referred to as inner-loop optimization, 2) Updating the model parameters by optimizing for the performance of f_{θ_i} with respect to θ across all sampled tasks performed as follows: $\theta \leftarrow \theta - \beta \nabla \sum_i L_{\tau_i}(f_{\theta_i})$, and referred to as outer-loop optimization.

Algorithm 1 Gradient Agreement Algorithm

- 1: Initialize model parameters, θ
 - 2: **for** iteration = 1, 2, ... **do**
 - 3: Sample a batch of N tasks $\tau_i \sim p(\tau)$
 - 4: **for** all τ_i **do**
 - 5: Split the examples of the task into k sub-batches
 - 6: $\theta_i = \theta - \alpha_{inner} \nabla L_{\tau_i}(f_\theta)$ for k steps
 - 7: $g_i = \theta - \theta_i$
 - 8: $w_i = \frac{\sum_{j \in T}(g_i^T g_j)}{\sum_{k \in T} |\sum_{j \in T}(g_k^T g_j)|}$ for all i
 - 9: $\theta = \theta - \alpha_{outer} \sum_i w_i \nabla L_{\tau_i}(f_{\theta_i})$ (MAML)
 - 10: $\theta = \theta + \alpha_{outer} \sum_i w_i (\theta_i - \theta)$ (Reptile)
 - 11: **return** θ
-

The assumption in the prior arts (Finn et al. [2017], Nichol et al. [2018]) is that each task should contribute equally to the direction of the outer-loop optimization step. However, in our approach, we associate a weight to the loss of each task, and the meta-optimization rule will become: $\theta \leftarrow \theta - \beta \nabla \sum w_i L_{\tau_i}(f_{\theta_i})$. Assuming that the gradient update vector for each task is presented by g_i , we define w_i as:

$$w_i = \frac{\sum_{j \in T}(g_i^T g_j)}{\sum_{k \in T} |\sum_{j \in T}(g_k^T g_j)|} \sim g_i^T g_{avg} \quad (1)$$

The proposed w_i is proportional to the inner product of the gradient of a task and the average of the gradients of all tasks in a batch. Therefore, if the gradient of a task is aligned with the average of the gradients of all tasks, it contributes more than other tasks to the next parameter update. With this insight, the proposed optimization method, pushes the model parameters towards initial parameters that more tasks agree upon. The full procedure of the proposed optimization is illustrated in Algorithm 1.

3 Analysis

In this section, inspired by (Jenni and Favaro [2018]), we provide a mathematical analysis for the proposed method. Denoting a batch of sampled tasks as T , the outer-loop update in batched version of MAML with current parameters θ^t will be:

$$\theta^{t+1} = \theta^t - \beta \sum_{i \in T} \nabla L_{\tau_i}(\theta_i) \quad (2)$$

Instead, in gradient agreement approach, we look for a linear combination of task losses that lead to a better approximation of the tasks errors. We use w_i per each task which is estimated at each iteration.

Therefore, the goal is to find parameters θ and coefficients w_i so that the model performs well on the sampled batch of tasks. We thus aim to optimize the following objective function:

$$\hat{\theta}, \hat{\mathbf{w}} = \arg \min_{\theta, \mathbf{w}} \sum_{i \in \mathbb{T}} L_{\tau_i}(\theta(\mathbf{w})) + \frac{\mu}{2} \|\mathbf{w}\|_2^2 \quad (3)$$

$$\text{subj. to: } \theta(\mathbf{w}) = \arg \min_{\bar{\theta}} \sum_{i \in \mathbb{T}} w_i L_{\tau_i}(\bar{\theta}), \quad \|\mathbf{w}\|_1 = 1 \quad (4)$$

where, \mathbf{w} is the vector of all weights, w_i and $i \in \mathbb{T}$, associated with each task, and μ is a regularization parameter for the distribution of weights. In this section, we show that the introduced inner product metric in 1 is corresponding to minimizing the objective function of Equation 3. Note that the solution of Equation 3 does not change if we multiply all the weights, w_i , by the same positive constant. As a result, we normalize the magnitude of \mathbf{w} to one. A typical method to solve the above optimization is the so-called *implicit differentiation* (Domke [2012]) method to solve a linear system in the second order derivatives of loss functions, which leads to solving a very high-dimensional linear system. To avoid computational complexities of the aforementioned method, we apply proximal approximation using Taylor series.

The loss function of i -th task at θ^t can be approximated using Taylor series as:

$$L_{\tau_i}(\theta) \approx L_{\tau_i}(\theta^t) + g_i(\theta - \theta^t) \quad (5)$$

where g_i is the update vector obtained in the inner-loop optimization for i -th task. By plugging in the first-order approximation of (5) to (3) and (4), we obtain the following objective function:

$$\theta^{t+1}, \hat{\mathbf{w}} = \arg \min_{\theta, \mathbf{w}} \sum_{i \in \mathbb{T}} [L_{\tau_i}(\theta^t) + g_i(\theta(\mathbf{w}) - \theta^t)] + \frac{\mu}{2} \|\mathbf{w}\|_2^2 \quad (6)$$

$$\text{subj. to: } \theta(\mathbf{w}) = \arg \min_{\bar{\theta}} \sum_{i \in \mathbb{T}} w_i [L_{\tau_i}(\theta^t) + g_i(\bar{\theta} - \theta^t)] + \frac{\|\bar{\theta} - \theta^t\|_2^2}{2\epsilon}, \quad \|\mathbf{w}\|_1 = 1 \quad (7)$$

The closed-form solution to the quadratic problem in (7) is a classical SGD update rule of:

$$\theta(\mathbf{w}) = \theta^t - \epsilon \sum_{i \in \mathbb{T}} w_i g_i \quad (8)$$

Now, we plug in the solution in (8) to (6):

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \sum_{i \in \mathbb{T}} \sum_{j \in \mathbb{T}} (-\epsilon w_i g_i^T g_j) + \frac{\mu}{2} \|\mathbf{w}\|_2^2 \quad (9)$$

We compute the derivative of (9) with respect to w_i for all $i \in T$ and set it to zero, temporarily ignoring the normalization constraint on \mathbf{w} :

$$\sum_{j \in \mathbb{T}} (-\epsilon g_i^T g_j) + \mu w_i = 0 \quad (10)$$

Therefore, w_i will be:

$$w_i = \lambda \sum_{j \in \mathbb{T}} (g_i^T g_j) \quad (11)$$

Now, by applying the L_1 constraint on the norm of \mathbf{w} , we obtain λ as:

$$\sum_{i \in \mathbb{T}} |w_i| = 1 \Rightarrow \lambda \sum_{i \in \mathbb{T}} \left| \sum_{j \in \mathbb{T}} (g_i^T g_j) \right| = 1 \Rightarrow \lambda = \frac{1}{\sum_{i \in \mathbb{T}} \left| \sum_{j \in \mathbb{T}} (g_i^T g_j) \right|} \quad (12)$$

As a result, w_i will be:

$$w_i = \frac{\sum_{j \in \mathbb{T}} (g_i^T g_j)}{\sum_{k \in \mathbb{T}} \left| \sum_{j \in \mathbb{T}} (g_k^T g_j) \right|} \quad (13)$$

4 Experiments and Results

4.1 One-Dimensional Sine Wave Regression

As a simple case study, consider the 1D sine wave regression problem. A task $\tau = (a, \phi)$ is defined by the amplitude a and phase ϕ of a sine wave function $f_\tau(x) = a\sin(x + b)$, a and ϕ are sampled from $U([0.1, 5.0])$ and $U([0, 2\pi])$, respectively. The average of L_2 loss over 1000 unseen sine wave tasks is 0.13 and 0.08 for Reptile and gradient agreement method, respectively.




Figure 1: Demonstration of Reptile and our method on a toy few-shot regression problem, where we train on 10 sampled points of a sine wave, performing 5 gradient steps on a MLP with layers $1 \rightarrow 64 \rightarrow 64 \rightarrow 1$. left: before training, middle: after Reptile training, right: after gradient agreement based training.

4.2 Few-shot Classification

We evaluate our method on two popular few-shot classification tasks, miniImageNet (Vinyals et al. [2016]) and Omniglot (Lake et al. [2011]). The Omniglot dataset consists of 20 instances of 1623 characters from 50 different alphabets. The miniImageNet dataset includes 64 training classes, 16 validation classes, and 20 test classes, each class contains 600 examples. The problem of N-way classification is set up as follows: we sample N classes from the total C classes in the dataset and then selecting K examples for each class. For our experiments, we use the same CNN architectures and data pre-processing as (Finn et al. [2017]) and (Nichol et al. [2018]). The results for miniImageNet and Omniglot are presented in Tables 1 and 2, respectively. The models learned by our approach compares well to the state-of-the-art results on these two tasks. It substantially outperforms the prior methods on the challenging task of miniImageNet.

Table 1: Accuracy results on miniImageNet.

Algorithm	5way 1shot	5way 5shot
MAML (Finn et al. [2017])	48.70±1.84%	63.11±0.92%
Reptile (Nichol et al. [2018])	49.97±0.32%	65.99±0.58%
Gradient Agreement	54.80±1.19%	73.27±0.62%

Table 2: Accuracy results on Omniglot.

Algorithm	5way 1shot	5way 5shot	20way 1shot	20way 5shot
MAML (Finn et al. [2017])	98.7±0.4%	99.9±0.1%	95.8±0.3%	98.9±0.2%
Reptile (Nichol et al. [2018])	97.68±0.04%	99.48±0.06%	89.43±0.14%	97.12±0.32%
Gradient Agreement	98.6±0.63%	99.8±0.12%	94.25±0.51%	98.99±1.34%

5 Conclusion and Future Works

In this work, we presented a generalization method for meta-learning algorithms that adjusts the model parameters by introducing a set of weights over the loss functions of tasks in a batch in order to maximize the dot products between the gradients of different tasks. The higher the inner products between the gradients of different tasks, the more agreement they have upon the model parameters update. We also presented the objective function of this optimization method by a theoretical analysis using first-order Taylor series approximation. This geometrical interpretation of optimization in meta-learning studies can be an interesting future direction.

References

- Ruslan Salakhutdinov, Brenden M. Lake, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. In *Science*, 2015.
- Jürgen Schmidhuber. Evolutionary principles in self-referential learning. Master’s thesis, 1987.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *CoRR*, abs/1703.03400, 2017. URL <http://arxiv.org/abs/1703.03400>.
- Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *CoRR*, abs/1803.02999, 2018. URL <http://arxiv.org/abs/1803.02999>.
- Oriol Vinyals, Charles Blundell, Timothy P. Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. Matching networks for one shot learning. *CoRR*, abs/1606.04080, 2016. URL <http://arxiv.org/abs/1606.04080>.
- Brenden M. Lake, Ruslan Salakhutdinov, Jason Gross, and Joshua B. Tenenbaum. One shot learning of simple visual concepts. In *CogSci*, 2011.
- Simon Jenni and Paolo Favaro. Deep bilevel learning. volume abs/1809.01465, 2018.
- Justin Domke. Generic methods for optimization-based modeling. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, 2012.