
Generalization of Model-Agnostic Meta-Learning Algorithms: Recurring and Unseen Tasks

Alireza Fallah
EECS Department
Massachusetts Institute of Technology
afallah@mit.edu

Aryan Mokhtari
ECE Department
The University of Texas at Austin
mokhtari@austin.utexas.edu

Asuman Ozdaglar
EECS Department
Massachusetts Institute of Technology
asuman@mit.edu

Abstract

In this paper, we study the generalization properties of Model-Agnostic Meta-Learning (MAML) algorithms for supervised learning problems. We focus on the setting in which we train the MAML model over m tasks, each with n data points, and characterize its generalization error from two points of view: First, we assume the new task at test time is one of the training tasks, and we show that, for strongly convex objective functions, the expected excess population loss is bounded by $\mathcal{O}(1/mn)$. Second, we consider the MAML algorithm’s generalization to an unseen task and show that the resulting generalization error depends on the total variation distance between the underlying distributions of the new task and the tasks observed during the training process. Our proof techniques rely on the connections between algorithmic stability and generalization bounds of algorithms. In particular, we propose a new definition of stability for meta-learning algorithms, which allows us to capture the role of both the number of tasks m and number of samples per task n on the generalization error of MAML.

1 Introduction

In several machine learning problems, it is of interest to design algorithms that can be adjusted based on previous experiences and tasks to perform better on a new task. In particular, meta-learning algorithms achieve such a goal through various approaches, including finding a proper meta-initialization for the new task [1–3], updating the model architecture [4–6], or learning the parameters of optimization algorithms [7, 8].

A popular meta-learning framework that has shown promise in practice is Model-Agnostic Meta-Learning (MAML), which was first introduced in [1]. MAML algorithm uses available training data on a number of tasks to come up with a meta-initialization that performs well after it is slightly updated at test time with respect to the new task. In other words, unlike standard supervised learning, in which we aim to find a model that generalize well to a new task *without any adaptation step*, in MAML our goal is to find an initial model for learning a new task when *we have access to limited labeled data for that task* to run one (or a few) step(s) of stochastic gradient descent (SGD).

As shown in Fig. 1, in MAML we are given m tasks with m corresponding datasets $\{\mathcal{S}_i\}_{i=1}^m$ in the training phase. Once the model is trained (w_{train}^*), a new task is revealed at test time for which we have access to K labeled samples drawn from $\mathcal{D}_{\text{test}}$. We use these labeled samples of the new task to

update the trained model by running a step of SGD leading to a new model for the test task (w_{new}^*). We finally evaluate the performance of the updated model over the test task, denoted by $\mathcal{L}_{\text{test}}(w_{\text{new}}^*)$.

MAML and its variants have been extensively studied over the past few years from both empirical and theoretical point of view [2, 9–16]. In particular, [13] provided convergence guarantees for MAML algorithm under the assumption that access to fresh samples at any round of the training stage is possible, and [15] extended this results to the case that multiple gradient steps can be performed at test time. However, one shortcoming of such analysis is that, at training stage, we often do not have access to fresh samples at every iteration. Instead, we have access to a large set of realized samples and we typically do multiple passes over the data points during the training stage.

Hence, it is essential to come up with a novel analysis that addresses this issue by characterizing the training error and generalization error of MAML separately. In this paper, we accomplish this goal and showcase the role of different problem parameters in the generalization error of MAML. Specifically, we assume that we are given m supervised learning tasks, with (possibly different) underlying distributions p_1, \dots, p_m , where for each task we have access to n samples¹. As we measure the performance of a model by its loss after one step of SGD adaptation with K samples, the problem that one can solve in the training phase is minimizing the average loss, over all given m tasks and their n samples, after one step of SGD with K samples. This empirical loss can be considered as a surrogate for the desired expected loss (with respect to tasks data) over all m tasks. Here, we focus on the case that MAML is used to solve this empirical minimization problem, and our goal is to quantify the test error of MAML output. To tackle this problem, we first briefly revisit the results from the optimization literature to bound the training error of MAML, assuming that the loss functions are strongly convex. We next turn to the main focus of our paper which is the generalization properties of MAML. More specifically, we address the following questions:

- *If one of the m given tasks recurs uniformly at random at test time, then how well (in expectation) would the trained model perform after adaptation with SGD over the fresh samples of that task?* In other words, having training error minimized, what would be the generalization error and our guarantee on test error? Here, we show that for strongly convex objective functions, we could achieve a generalization error that decays at $\mathcal{O}(1/mn)$. Our analysis builds on the connections between algorithmic stability and generalization of the output of algorithms. While this relation is well-understood in classic statistical learning [17, 18], here we propose a novel stability definition for meta-learning algorithms which allows us to restore such connection for our setting.
- *Assuming that the task at test time is NOT one of the m tasks at training, how would the model perform on that task after the adaptation step?* We answer this question by focusing on the case that the revealed task at the test time is a new *unseen* task with underlying data distribution p_{m+1} , and formally characterizing the generalization error of MAML in this case. We show that when the task at test time is new, the generalization error also depends on the total variation distance between p_{m+1} and p_1, \dots, p_m .

Related work: Recently, there has been significant progress in studying theoretical aspects of meta-learning, in particular, MAML. Authors in [19] proposed iMAML which updates the model using an approximation of one step of proximal point method and studied its convergence. In [20], authors introduced the task-robust MAML by considering a minimax formulation rather than minimization. Several papers have also studied MAML through more general frameworks such as bilevel optimization [21], stochastic compositional optimization [22], and conditional stochastic

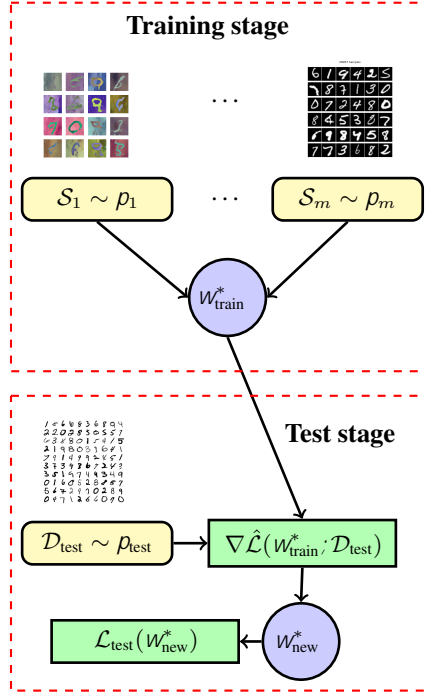


Figure 1: MAML framework

¹More precisely, in our analysis we take $2n$ samples per each task to simplify derivations.

optimization [23]. Also, several works have studied the extension of meta-learning theory to online learning [24, 3], federated learning [25], and reinforcement learning [26, 27].

The most relevant paper to our work is [28] that studies generalization of meta-learning algorithms using stability techniques and shows a $\mathcal{O}(1/\sqrt{m})$ bound for nonconvex loss functions. Here we focus on strongly convex objective functions and present an analysis that differs from this work in two fundamental aspects. First, we present a different notion of stability that allows us to capture the number of data points per task in our bound. In particular, our stability notion measures sensitivity of the algorithm to perturbations that involve changing K data points which is the data unit involved in the adaptation step of the MAML algorithm. This enables us to obtain a much tighter bound $\mathcal{O}(1/mn)$ (compared to $\mathcal{O}(1/m)$ achieved in [28] for strongly convex functions), highlighting the dependence on the number of the data samples available for each task. Second, we also consider the generalization of MAML for the case that the task at test time is not one of the available tasks during the training stage.

The generalization of MAML has also been studied in [29] from an empirical point of view. In particular, they show that the generalization of MAML to new tasks is correlated with the coherence between their adaptation trajectories in parameter space. This is aligned with the connection of generalization and closeness of underlying distributions that we observe in our results.

2 Problem formulation

In this paper, we consider the supervised learning setting, where each data point is denoted by $z = (x, y) \in \mathcal{Z}$ with $x \in \mathcal{X}$ being the input (feature vector) and $y \in \mathcal{Y}$ being its corresponding label. We use the loss function $l : \mathcal{R}^d \times \mathcal{Z} \rightarrow \mathcal{R}^+$ to evaluate the performance of a model parameterized by $w \in \mathcal{W}$, where \mathcal{W} is a convex and closed subset of \mathcal{R}^d . In other words, for a data point $z = (x, y) \in \mathcal{Z}$, the loss $\ell(w, z)$ denotes the error of model w in predicting the label y given input x .

We consider access to m tasks denoted by $\mathcal{T}_1, \dots, \mathcal{T}_m$, where the data corresponding to each task \mathcal{T}_i is generated from a distinct distribution p_i . The *population loss* corresponding to task \mathcal{T}_i for model w is defined as $\mathcal{L}_i(w) := \mathbb{E}_{z \sim p_i}[\ell(w, z)]$.

We further use the notation $\hat{\mathcal{L}}(w; \mathcal{D})$ to denote the *empirical loss* corresponding to dataset \mathcal{D} , which is defined as the average loss of w over the samples of dataset \mathcal{D} , i.e., $\hat{\mathcal{L}}(w; \mathcal{D}) := \frac{1}{|\mathcal{D}|} \sum_{z \in \mathcal{D}} \ell(w, z)$, where $|\mathcal{D}|$ is the size of dataset \mathcal{D} . In general, and throughout the paper, we use the *hat notation* to distinguish *empirical* losses from *population* losses.

Our goal is to find $w \in \mathcal{W}$ that performs well on average² over all tasks, after it is updated with respect to the new task and by using one step of stochastic gradient descent (SGD) with a batch of size K . To formally introduce this problem we first define the function $F_i(w)$ which captures the performance of model w over task \mathcal{T}_i once it is updated by a single step of SGD,

$$F_i(w) := \mathbb{E}_{\mathcal{D}_i^{\text{test}}} \mathcal{L}_i(w - \alpha \nabla \hat{\mathcal{L}}(w, \mathcal{D}_i^{\text{test}})) = \mathbb{E}_{\mathcal{D}_i^{\text{test}}} \mathbb{E}_{z \sim p_i} \ell(w - \frac{\alpha}{K} \sum_{z^0 \in \mathcal{D}_i^{\text{test}}} \nabla \ell(w, z'), z) \quad (1)$$

where $\mathcal{D}_i^{\text{test}}$ is a batch with K different samples, drawn from the probability distribution p_i . Note that the outer expectation is taken with respect to the choice of elements of $\mathcal{D}_i^{\text{test}}$ while the inner one is taken with respect to the data of task i .

As our goal is to find a model that performs well after one step of adaptation over all m tasks, we minimize the average expected loss over all given tasks, which can be written as

$$\min_{w \in \mathcal{W}} F(w) := \frac{1}{m} \sum_{i=1}^m F_i(w). \quad (2)$$

As the underlying distribution of tasks are often unknown in most applications, we are often unable to directly solve the problem in (2). On the other hand, for each task, we often have access to data points that are drawn according to their data distributions. Therefore, instead of solving (2), we solve its sample average surrogate problem in which each F_i is approximated by its empirical loss.

²Our analysis can be extended to the case that the distribution over tasks is not uniform.

To formally define the empirical loss for each task, suppose for each task \$i\$ we have access to a training set \$S_i\$, where its elements are drawn independently according to the probability distribution \$p_i\$. We further divide the set \$S_i\$ into two disjoint sets of size defined as \$S_i^{in}\$ and \$S_i^{out}\$, i.e., \$S_i := S_i^{in} \cup S_i^{out}\$ and \$|S_i^{in}| = |S_i^{out}| = n\$. Here, we use the elements of \$S_i^{in}\$ to estimate the inner gradient \$\nabla_w \hat{L}(w; D_i^{in})\$ and use the samples in the set \$S_i^{out}\$ to estimate the outer function \$f_i(\cdot)\$. Specifically, we define the sample average \$\hat{F}_i\$ using data sets \$S_i^{in}\$ and \$S_i^{out}\$ as

$$\begin{aligned} \hat{F}_i(w; S_i) &:= \frac{1}{n} \sum_{\substack{D_i^{in} \subseteq S_i \\ |D_i^{in}|=K}} \hat{L}(w; D_i^{in}) \quad \nabla_w \hat{L}(w; D_i^{in}); S_i^{out} \quad (3) \\ &= \frac{1}{n} \sum_{\substack{D_i^{in} \subseteq S_i \\ |D_i^{in}|=K}} \frac{1}{K} \sum_{z \in S_i^{out}} \nabla_w \hat{L}(w; z^0); z^A \end{aligned}$$

This expression shows that we use elements of \$S_i^{out}\$ to approximate the expectation required for the computation of \$F_i\$, and we approximate the expectation with respect to the test set by averaging over all subsets of \$S_i^{in}\$ that have \$K\$ elements. Given this expression, the sample average approximation (empirical loss) of Problem (1) is given by

$$\arg \min_{w \in \mathcal{W}} \hat{F}(w; S) := \frac{1}{m} \sum_{i=1}^m \hat{F}_i(w; S_i) \quad (4)$$

where \$S := \{S_i\}_{i=1}^m\$ is defined as the concatenation of all tasks data sets.

Having the dataset \$S\$, a (possibly randomized) optimization algorithm \$\mathcal{A}\$ with output \$\mathcal{A}(S)\$ can be used to find an approximate solution to the problem (4). The error of this solution with respect to the MAML empirical loss, i.e. \$\hat{F}(\mathcal{A}(S); S) - \min_w \hat{F}(\cdot; S)\$, is called training error. In this paper, we are mainly interested to bound the test error which is the error of \$\mathcal{A}(S)\$ with respect to the population loss, i.e. \$F(\mathcal{A}(S)) - \min_w F\$. The test error is also sometimes called excess (population) loss. Note that the expected test error can be decomposed into three terms:

$$\begin{aligned} E_{A;S} [F(\mathcal{A}(S)) - \min_w F] &= \text{(test error)} \\ &= \underbrace{E_{A;S} [F(\mathcal{A}(S)) - \hat{F}(\mathcal{A}(S); S)]}_{\text{generalization error}} + \underbrace{E_{A;S} [\hat{F}(\mathcal{A}(S); S) - \min_w \hat{F}(\cdot; S)]}_{\text{training error}} + \underbrace{E_S [\min_w \hat{F}(\cdot; S) - \min_w F]}_{z_0} \end{aligned}$$

It can be verified that the expectation of the third term (over \$S\$) is non-positive since \$E_S[\min_w \hat{F}(\cdot; S)] \leq \min_w E_S[\hat{F}(\cdot; S)]\$ and \$E_S[\hat{F}(\cdot; S)] = F\$. Hence, to bound the expected test error, we should bound the expectation of training and generalization errors.

The Model-Agnostic Meta-Learning (MAML) method proposed in [1] is designed to solve the empirical minimization problem defined in (4). The steps of MAML are outlined in Algorithm 1. MAML solves Problem (4) by using SGD update for the average loss function \$\hat{F}(w; S)\$. To better highlight this point, note that the gradient of \$\hat{F}(w; S)\$ can be written as \$\frac{1}{m} \sum_{i=1}^m \nabla_w \hat{F}_i(w; S_i)\$, where the \$i\$-th term corresponding to task \$i\$ is given by

$$\nabla_w \hat{F}_i(w; S_i) = \frac{1}{n} \sum_{\substack{D_i^{in} \subseteq S_i \\ |D_i^{in}|=K}} (I_d - \nabla_w^2 \hat{L}(w; D_i^{in})) \nabla_w \hat{L}(w; D_i^{in}); S_i^{out} \quad ; \quad (5)$$

which involves the second-order information of the loss function. Therefore, to compute a mini-batch approximation for the above gradient, we consider the batch sets \$S_i^{in}\$ with size \$K\$ and \$D_i^{out} \subseteq S_i^{out}\$ with \$b\$ elements. Replacing the above sums with their batch approximations leads to the following stochastic gradient approximation

$$g_i(w; D_i^{in}; D_i^{out}) := (I_d - \nabla_w^2 \hat{L}(w; D_i^{in})) \nabla_w \hat{L}(w; D_i^{in}); D_i^{out} \quad ; \quad (6)$$

which is indeed an unbiased estimator of the gradient \$\nabla_w \hat{F}_i(w; S_i)\$ in (5). If for each task we perform the update of SGD with \$g_i\$ and then compute their average it would be similar to running SGD for

Algorithm 1: MAML [1]

Input: The set of datasets $\mathcal{S} = \{S_i\}_{i=1}^m$ with $S_i = \{S_i^{\text{in}}, S_i^{\text{out}}\}$; test time batch size k ; # of tasks summoned at each round $\#$; of iterations T .
 Choose arbitrary initial point $w^0 \in \mathbb{R}^d$;
 for $t = 0$ to $T - 1$ do
 Choose $\#$ tasks uniformly at random (out of m tasks) and store their indices B_t ;
 for all T_i with $i \in B_t$ do
 Sample a batch $D_i^{t,\text{in}}$ of K different elements from S_i^{in} with replacement;
 Sample a batch $D_i^{t,\text{out}}$ of size b from S_i^{out} and with replacement;
 $w_i^{t+1} := w^t - \eta \nabla L(w^t; D_i^{t,\text{in}}) - \eta \nabla L(w^t; D_i^{t,\text{out}})$;
 end for
 $w^{t+1} := \frac{1}{\#} \sum_{i \in B_t} w_i^{t+1}$;
 end for
 Return: w^T and $\bar{w}^T := \frac{1}{T+1} \sum_{t=0}^T w^t$

the average loss $\mathbb{E} F(w; S)$. This is exactly how MAML is implemented in practice as outlined in Algorithm 1. In this paper, we consider a constrained problem, and as a result, we also need an extra projection step in the last step to ensure the feasibility of iterates. Finally, the output of MAML could be the last iterate w^T or the time-average of all iterates $\bar{w}^T := \frac{1}{T+1} \sum_{t=0}^T w^t$.

As stated earlier, the convergence properties of MAML-type methods from an optimization point of view have been studied recently under different set of assumptions. In this paper, as we characterize the sum of training error and generalization error, we briefly discuss the optimization error of MAML when it is used to solve the empirical problem (4). However, the main focus of this paper is on studying the generalization error of MAML with respect to new samples and new tasks. Specifically, we aim to address the following questions: How well does the solution of (4) generalize to the main problem of interest in (2)? This could be seen as the generalization error of the MAML algorithm over new samples for recurring tasks. How well does the solution of (4) generalize to samples from new unseen tasks? To be more precise, how would the obtained model perform if the new task is not one of the m tasks T_1, \dots, T_m observed at training, and it is rather a new unseen task T_{m+1} with an unknown underlying distribution p_{m+1} ? In the upcoming sections, we answer these questions on the generalization properties of MAML in detail and characterize the role of number of tasks, number of samples per task, and number of labeled samples revealed at test time.

3 Theoretical results

In this section, we formally characterize the excess population loss (test error) of the MAML solution, when we measure the performance of a model after one step of SGD adaptation. In particular, we first discuss the training error of MAML in detail. Then, we establish a generalization error bound for the case that the solution of MAML is evaluated over new samples of a recurring task. Finally, we state the generalization error of MAML once its solution is applied to a new unseen task. Before stating our results, we mention our required assumptions.

Assumption 1. For any $z \in \mathbb{Z}$, the function $\ell(\cdot; z)$ is twice continuously differentiable. Furthermore, we assume it satisfies the following properties for any $u \in \mathbb{R}^d$:

- (i) For any $z \in \mathbb{Z}$, the function $\ell(\cdot; z)$ is μ -strongly convex, i.e., $\ell(w; z) - \ell(u; z) \geq \frac{\mu}{2} \|w - u\|^2$;
- (ii) The gradient norm is uniformly bounded over \mathbb{W} , i.e., $\|\nabla \ell(w; z)\| \leq G$;
- (iii) The loss is L -smooth over \mathbb{R}^d , i.e., $\|\nabla \ell(w; z) - \nabla \ell(u; z)\| \leq L \|w - u\|$;
- (iv) Hessian is β -Lipschitz continuous over \mathbb{R}^d , i.e., $\|\nabla^2 \ell(w; z) - \nabla^2 \ell(u; z)\| \leq \beta \|w - u\|$.

We also require the following assumption on the tasks distribution. This assumption implies that, with probability one, a set of finite samples generated from a distribution are all different.

Assumption 2. We assume \mathcal{Z} is a Polish space (i.e., complete, separable, and metric) and \mathcal{F} is the Borel σ -algebra over \mathcal{Z} . Moreover, for any j , p_j is a non-atomic probability distribution over $(\mathcal{Z}; \mathcal{F}_{\mathcal{Z}})$, i.e., $p_j(z) = 0$ for every $z \in \mathcal{Z}$.

3.1 Training error

While the main focus of this paper is on studying the population error of MAML algorithm, we first study its training error which is required to provide characterization of the excess loss of MAML. To do so, we first state the following result from [24] and [13] on the strong convexity and smoothness of $\ell(w; r; \hat{L}(w; D); z)$ for any batch D and any $z \in \mathcal{Z}$.

Lemma 1 ([13] & [24]). If Assumption 1 holds, then for an arbitrary batch D and $z \in \mathcal{Z}$, and with $\frac{1}{L}$, the function $\ell(w; r; \hat{L}(w; D); z)$ is $4L + 2G$ smooth over W . Furthermore, $\ell(w; r; \hat{L}(w; D); z)$ is $\frac{1}{8}$ -strongly convex, if $\min\{\frac{1}{2L}; \frac{1}{8G}\}g$.

An immediate consequence of this Lemma is that the MAML empirical loss defined in (4) is also $\frac{1}{8}$ -strongly convex and $4L + 2G$ smooth over W . In addition, it can be shown that the norm of $g_i(w; D_i^{\text{in}}; D_i^{\text{out}})$ defined in (6), which is the unbiased gradient estimate used in MAML, is uniformly bounded above; for more details check Lemma 5 in Appendix A. Having these properties of the MAML empirical loss established, we next state the following proposition on the training error of MAML. This result is obtained by slightly modifying the well-known results on the convergence of SGD in [30–32] in order to take into account the stepsize constraints that are imposed by generalization analysis. For completeness, the proof of this result is provided in Appendix B.

Proposition 1. Consider $\hat{F}(\cdot; S)$ defined in (4) with $\min\{\frac{1}{2L}; \frac{1}{8G}\}g$. If Assumption 1 holds, then for MAML with $\eta = \min\{\frac{1}{L}; \frac{1}{8(t+1)}\}$ for $t = 0, \dots, T-1$, and for any $S \in \mathcal{S}$, the last iterate w^T satisfies

$$\mathbb{E} \left[\hat{F}(w^T; S) - \hat{F}(w_S; S) \right] \leq O(1) \frac{G^2(1 + \frac{1}{L})}{2} \frac{L + G}{T} + \frac{G}{T}; \quad (7)$$

and the time-average of iterates w^t satisfies

$$\mathbb{E} \left[\hat{F}(w^T; S) - \hat{F}(w_S; S) \right] \leq O(1) \frac{G^2(\log(T) + \frac{1}{L})}{T};$$

where $w_S := \arg \min_{w \in W} \hat{F}(\cdot; S)$ and the expectations are taken over the randomness of algorithm.

In the above expressions, the notation $O(1)$ only hides absolute constants. It is worth noting that the term $\frac{G}{T}$ in (7) vanishes, if w_S be a minimizer of the unconstrained problem, i.e., $\hat{F}(w_S; S) = 0$.

3.2 Generalization error

We derive our generalization bounds for MAML by establishing its algorithmic stability properties. The stability approach has been used widely to characterize the generalization properties for optimization algorithms such as stochastic gradient descent [16] or differentially private methods [33]. These arguments are based on showing the uniform stability of algorithms [17] which we restate it here.

Definition 1 ([17]). Consider the problem of minimizing the empirical function $\hat{L}(w; H)$ for some dataset H . A randomized algorithm \mathcal{A} with output w_H given dataset H is called ϵ -uniformly stable if the following condition holds: Take the dataset H' which is the same as H , except at one data points. Then, we have $\sup_{z \in \mathcal{Z}} \mathbb{E}_{\mathcal{A}} [|\ell(w_H; z) - \ell(w_{H'}; z)|] \leq \epsilon$; where the expectation is taken over the randomness of \mathcal{A} .

The above definition captures the stability of an algorithm. Specifically, it states that Algorithm \mathcal{A} is ϵ -stable, if the resulting loss of its outputs, when it is run using two different datasets that only differ in one data point, are at most ϵ away from each other. Note that the above definition holds if the difference between the losses evaluated at any point z is bounded by ϵ . The main importance of this definition is its connection with generalization error. In particular, it can be shown that if an algorithm is ϵ -uniformly stable and “symmetric”, then its generalization error is bounded above by ϵ , see, e.g., [17]. Next, we formally state the definition of a symmetric algorithm.

Definition 2. An algorithm $A : Z^n \rightarrow \mathbb{R}^d$ is called symmetric, if for any $S \in Z^n$, the distribution of its output, i.e. $A(S)$, does not depend on the ordering of elements of S , i.e., if we take S^0 as a permutation of S , the distribution of $A(S)$ and $A(S^0)$ would be similar.

Note that Definition 1 is useful for the case where we measure the performance of a model's loss function over a sample, i.e. $\ell(w; z)$. However, in this paper we measure the performance of a model by looking at its loss after one step of SGD which involves data points, as defined in (6). Therefore, we cannot directly use Definition 1 for characterizing the generalization error of MAML. In fact, in what follows, we first propose a modified version of the uniform stability definition, which is compatible with our setting, and then show how such stability could lead to generalization bounds for MAML-type algorithms.

Definition 3. Consider the problem in (4). A randomized algorithm A with output w_S given dataset S is called $(\epsilon; K)$ -uniformly stable if the following condition holds for any $g \in \{1, \dots, m\}$: Take the dataset S which is the same as S , except that S_i^{in} and S_i^{out} differ from S_i^{in} and S_i^{out} in at most K and one data points, respectively. Then, for any Z and any K distinct points $z_1, \dots, z_K \in Z$,

$$\mathbb{E}_A \left[\sum_{i=1}^m \ell(w_S; z_i) \right] - \mathbb{E}_A \left[\sum_{i=1}^m \ell(w_{S'}; z_i) \right] \leq \epsilon;$$

where the expectation is taken over the randomness of

A few remarks about the above definition follow. First, one might wonder, why it is needed to change K points of the set S_i^{in} , while we change only one point of the set S_i^{out} . Note that, going from (1) to (3), the expectation $\mathbb{E}_{D_i^{\text{test}}}$ is replaced by the sum over all possible batches S_i^{in} of size K from S_i^{in} . In other words, for the empirical sum (6), each batch D_i^{in} can be seen as a data unit. That said, and similar to Definition 1, to characterize the stability, we need to change one data unit which is one batch of size K . That is why we change K data points of S_i^{in} in the definition of $(\epsilon; K)$ -uniformly stable. On the other hand, we replace $\ell(\cdot) = \mathbb{E}_{z \sim p_i}[\ell(\cdot; z)]$ in (1) with a sum over points of S_i^{out} in (3), and thus, for this one, each data unit is just a single data point. So, similar to Definition 1, we just change one data point for the set S_i^{out} .

Second, it is worth comparing this definition with the other definition given for stability of meta-learning algorithms in [28]. In that paper, the definition of stability is based on modifying the whole dataset S_i rather than what we do here which is changing just 1 points. While taking such a definition makes the analysis relatively simpler, it prohibits us from characterizing the dependence of generalization error on n , and hence the resulting upper bound for generalization error would be larger. We will come back to this point later when we derive the stability of MAML with respect to Definition 3 and compare it with the one obtained in [28].

As we discussed, the main reason that we are interested in the uniform stability of an algorithm is its connection with generalization error. In the next theorem, we formalize this connection for MAML formulation and show that if an Algorithm A is $(\epsilon; K)$ -uniformly stable and symmetric, then its output generalization error is bounded above by ϵ . The proof of this result is available in Appendix C.

Theorem 1. Consider the population and empirical losses defined in (2) and (4), respectively. If Assumption 2 holds and A is a (possibly randomized) symmetric and $(\epsilon; K)$ -uniformly stable algorithm with output w_S , then $\mathbb{E}_{A; S} [F(w_S)] - F(w_S; S) \leq \epsilon$.

This result shows that if we prove a symmetric algorithm is $(\epsilon; K)$ -uniformly stable as defined in Definition 3, then we can bound its output model generalization error. Hence, to characterize the generalization error of the model trained by MAML algorithm, we only need to capture the uniform stability parameter of MAML. Before stating this result, it is worth noting that while we limit our focus to MAML in this paper, Definition 3 and Theorem 1 could provide a framework for studying the generalization properties of a broader class of gradient-based meta-learning algorithms such as Reptile [34], First-order MAML [1], and Hessian-Free MAML [13].

Theorem 2. If Assumption 1 holds, then MAML (Algorithm 1) with both last iterate and average iterate outputs and with $\eta = \frac{1}{2L} \frac{1}{8\bar{G}}$ and $\tau = \frac{1}{4L+2} \frac{1}{\bar{G}}$ is $(\epsilon; K)$ -uniformly stable, where

$$\epsilon := O(1) \frac{G^2(1+LK)}{mn};$$

According to the above discussion, the result of Theorem 2 guarantees that the generalization error of MAML solution decays by a factor $O(1/(K=mn))$, where m is the number of tasks in the training set

and n is the number of available samples per task. The classic lower bound for SGD over strongly convex functions translates to a $(1=mn)$ lower bound in our setting. Hence, our bound is tight in the small K regime, which is generally the case in few-shot learning problems. However, one shortcoming of this result is that it is not tight in the large regime. In Appendix E we show how we could improve this result for the large regime. However, throughout the paper, we keep our discussion limited to the small regime.

Remark 1. If instead of using our uniform stability definition (i.e., Definition 3), one uses the stability definition given in [28], the resulted stability constant would be proportional to $(1=mn)$ rather than $(1=mn)$. In fact, our proposed uniform-stability definition empowers us to obtain a better bound and indicates the role of number of samples per task on the generalization error.

Remark 2. The algorithmic stability technique is mainly limited to the convex setting, since, in the nonconvex case, we need to keep learning rate very small to obtain meaningful generalization results which makes it impractical (Check Appendix G for further discussions on this matter). In fact, the main reason that we assume strongly convex and (\cdot) is to ensure that the meta-objective is convex, as, in general, relaxing any of these two could lead to a nonconvex meta-objective function. However, these two assumptions together make the objective function strongly convex, which is not necessarily needed in our analysis. In fact, if we assume that the meta-function are convex (but not necessarily strongly convex), we could still use Definition 3 to derive similar generalization bounds.

Putting Proposition 1 and Theorem 2 together, we obtain the following result on the excess population loss of MAML algorithm. We only report the result for the averaged iterates here, but one can obtain the result for the last iterate similarly by using Proposition 1.

Proposition 2. Consider the function F defined in (2) with $\min_{\mathcal{W}} \frac{1}{2L} ; \frac{1}{8G} g$. If Assumptions 1 and 2 hold, then the average of iterates generated by MAML (Algorithm 1) with $\min(\frac{1}{4L+2} \frac{1}{G} ; \frac{8}{(t+1)})$ after T iterations satisfies

$$\mathbb{E}_{A;S} \mathbb{E}_{\mathcal{W}} F(w^T) - \min_{\mathcal{W}} F(w^T) = O(1) \frac{G^2}{T} \frac{\log(T) + L}{T} + \frac{1 + LK}{mn} ;$$

where the expectation is taken over the sampling and the randomness of MAML algorithm.

As an immediate application, the following corollary characterizes MAML test error.

Corollary 1. Under the premise of Proposition 2, MAML algorithm after $T = O(mnL)$ iterations returns w^T such that $\mathbb{E}_{A;S} F(w^T) - \min_{\mathcal{W}} F(w^T) = O(G^2(1 + LK) / (mn))$:

3.3 Generalization to an unseen task

As we discussed in Section 2, another generalization measure is how the model trained with respect to the empirical problem in (4) performs on a new and unseen task k_1 with corresponding distribution p_{m+1} . To state our result for this case, we first need to introduce the following distance notion between probability distributions.

Definition 4. For two distributions P and Q , defined over the sample space and \mathcal{F} -field F , the total variation distance is defined as $k_{TV}(P, Q) := \sup_{A \in \mathcal{F}} |P(A) - Q(A)|$;

It is well-known that the total variation distance admits the following characterization

$$k_{TV}(P, Q) = \sup_{f: \mathcal{X} \rightarrow \{0, 1\}} \mathbb{E}_x \sim P[f(x)] - \mathbb{E}_x \sim Q[f(x)]; \quad (8)$$

Also, we require the following boundedness assumption for our result.

Assumption 3. For any $z \in \mathcal{Z}$, the function $\psi(\cdot; z)$ is M -bounded over \mathcal{W} .

Considering these assumptions, we are ready to state our result for the case when the task at test time is a new task and is not observed during training.

Theorem 3. Consider the population losses defined in (1) and (2). Suppose Assumptions 1, 2 and 3 hold. Then, for any $w \in \mathcal{W}$, we have

$$|\mathbb{E}_{p_{m+1}} F(w) - F(w)| \leq D(p_{m+1}; p_{g_{i=1}^m}); \quad (9)$$

where

$$D(p_{m+1}; f_{p_i g_{i=1}^m}) := \frac{4G^2 \chi^n}{m} \sum_{i=1}^m k_{p_{m+1}} p_i k_{TV} + (M + 2G^2) k_{p_{m+1}} \frac{1}{m} \sum_{i=1}^m p_i k_{TV} : (10)$$

While the proof is provided in detail in Appendix F, here we discuss a sketch of it to highlight the main technical contributions. To simplify the notation here, let us assume p_1 is the distribution used for training and p_2 is the distribution corresponding to the new task. Note that we aim to bound $|F_2(w) - F_1(w)|$. Recalling the definition of population loss (9), we need to bound the following expression (we drop the absolute value due to symmetry)

$$E_{f_{z_j^2} p_2 g_{j=1}^K; z^2} |w - r \hat{L}(w; f_{z_j^2} g_j); z^2| - E_{f_{z_j^1} p_1 g_{j=1}^K; z^1} |w - r \hat{L}(w; f_{z_j^1} g_j); z^1| : (11)$$

Notice that this difference can be cast as $E_{(f_{z_j} g_{j=1}^K; z)} p_2^{K+1} [X] - E_{(f_{z_j} g_{j=1}^K; z)} p_1^{K+1} [X]$; with

$X := |w - r \hat{L}(w; f_{z_j} g_{j=1}^K); z|$. As a result, a naive approach would be using Lipschitz and boundedness properties (Assumptions 1 and 3) along with (6) to obtain a bound depending on $k_{p_1}^{K+1} p_2^{K+1} k_{TV} = O(K) k_{p_1} p_2 k_{TV}$. However, this bound is not tight as it grows with

To address this issue, we exploit a coupling technique. Note that the expression (11) does not depend on the joint distribution of z_j^1 and z_j^2 , and instead, it only depends on the marginal distribution of z_j^1 and z_j^2 . That said, for each j , we assume that z_j^1 and z_j^2 are sampled from a distribution on $Z \times Z$ such that $z_j^1 \sim p_1, z_j^2 \sim p_2$, and $(z_j^1 \in Z^2) = k_{p_1} p_2 k_{TV}$. Such a coupling exists and is called maximal coupling of p_1 and p_2 [35]. Using this idea, as we show in Appendix F, we can eliminate the dependence on n , and as a result, the upper bound in (9) is independent of number of available labeled samples at test time denoted by

Remark 3. Note that the terms $\frac{1}{m} \sum_{i=1}^m k_{p_{m+1}} p_i k_{TV}$ and $k_{p_{m+1}} \frac{1}{m} \sum_{i=1}^m p_i k_{TV}$ in $D(p_{m+1}; f_{p_i g_{i=1}^m})$ come from the fact that we consider uniform distribution over tasks in the empirical problem (4). In particular, if we instead consider the empirical problem $\arg \min_{w \in \mathcal{W}} \sum_{i=1}^m q_i F_i(w; S_i)$; for some non-negative weights q_i with $\sum_{i=1}^m q_i = 1$, then $D(p_{m+1}; f_{p_i g_{i=1}^m})$ on the right hand side of (9) would change to

$$(M + 2G^2) k_{p_{m+1}} \sum_{i=1}^m q_i p_i k_{TV} + 12G^2 \sum_{i=1}^m q_i k_{p_{m+1}} p_i k_{TV} :$$

This result shows that by changing the training problem we can achieve a lower generalization error for MAML, if we have some information about the distribution p_{m+1} at training time. For instance, if we know p_{m+1} will be much closer to p_1 compared to p_2 , making the weight q_1 larger than q_2 would decrease the generalization error of MAML.

Corollary 2. Recall the population loss F_{m+1} defined in (2) and $D(p_{m+1}; f_{p_i g_{i=1}^m})$ defined in Theorem 3. Let A be an algorithm for solving the empirical problem (4) which achieves excess risk, i.e., $E_{A; S} [F(A(S))] = \min_{\mathcal{W}} F$. If Assumptions 1, 2 and 3 hold, then algorithm A finds a model w_S which achieves $+ D(p_{m+1}; f_{p_i g_{i=1}^m})$ excess loss with respect to p_{m+1} ,

$$E_{A; S} [F_{m+1}(w_S)] = \min_{\mathcal{W}} F_{m+1} + 2D(p_{m+1}; f_{p_i g_{i=1}^m}) :$$

This corollary and Proposition 2 together imply that the MAML algorithm's test error with respect to the new task p_{m+1} is $O(1) \frac{1}{mn} + D(p_{m+1}; f_{p_i g_{i=1}^m})$. As a result, if the new task's distribution p_{m+1} is sufficiently close to the other tasks' distributions, MAML will have a low test error on the new unseen task. On the other hand, if p_{m+1} is far from p_1, \dots, p_m in TV distance, then test error of the model trained over $T_{i=1}^m g$ over T_{m+1} could be potentially large. In Appendix F.2 we show how this result can be extended to the case that the task at test time is generated from a distribution over both recurring tasks $T_{i=1}^m g_{i=1}^m$ and the unseen task T_{m+1} .

4 Conclusion and future work

In this work, we studied the generalization of MAML algorithm in two key cases when the test time task is a recurring task from the ones observed during the training stage, or it is a

new and unseen one. For the first one, and under strong convexity assumption, we showed that the generalization error improves as the number of tasks or the number of samples per task increases. For the second case, we showed that when the distance between the unseen task's distribution and the distributions of training tasks is sufficiently small, the MAML output generalizes well to the new task revealed at test time.

While we focused on the convex case in this paper, deriving generalization bounds when the meta-function is nonconvex is a natural future direction to explore. However, this could be challenging since the generalization of gradient methods is not well understood in the nonconvex setting even for the classic supervised learning problem.

5 Acknowledgment

Alireza Fallah acknowledges support from the Apple Scholars in AI/ML PhD fellowship and the MathWorks Engineering Fellowship. This research is sponsored by the United States Air Force Research Laboratory and the United States Air Force Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. This research of Aryan Mokhtari is supported in part by NSF Grant 2007668, ARO Grant W911NF2110226, the Machine Learning Laboratory at UT Austin, and the NSF AI Institute for Foundations of Machine Learning.

References

- [1] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in Proceedings of the 34th International Conference on Machine Learning (Sydney, Australia), 06–11 Aug 2017.
- [2] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," arXiv preprint arXiv:1803.02999, 2018.
- [3] M. Khodak, M.-F. F. Balcan, and A. S. Talwalkar, "Adaptive gradient-based meta-learning methods," in Advances in Neural Information Processing Systems, pp. 5915–5926, 2019.
- [4] B. Baker, O. Gupta, N. Naik, and R. Raskar, "Designing neural network architectures using reinforcement learning," in International Conference on Learning Representations, 2017.
- [5] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," in International Conference on Learning Representations, 2017.
- [6] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 8697–8710, 2018.
- [7] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," in International Conference on Learning Representations, 2017.
- [8] M. Andrychowicz, M. Denil, S. Gómez, M. W. Hoffman, D. Pfau, T. Schaul, B. Shillingford, and N. de Freitas, "Learning to learn by gradient descent by gradient descent," in Advances in Neural Information Processing Systems, pp. 3981–3989, Curran Associates, Inc., 2016.
- [9] A. Antoniou, H. Edwards, and A. Storkey, "How to train your MAML," in International Conference on Learning Representations, 2019.
- [10] Z. Li, F. Zhou, F. Chen, and H. Li, "Meta-SGD: Learning to learn quickly for few-shot learning," arXiv preprint arXiv:1707.09835, 2017.
- [11] E. Grant, C. Finn, S. Levine, T. Darrell, and T. Griffiths, "Recasting gradient-based meta-learning as hierarchical bayes," International Conference on Learning Representations 2018.

- [12] H. S. Behl, A. G. Baydin, and P. H. S. Torr, “Alpha MAML: adaptive model-agnostic meta-learning,” 2019.
- [13] A. Fallah, A. Mokhtari, and A. Ozdaglar, “On the convergence theory of gradient-based model-agnostic meta-learning algorithms,” *International Conference on Artificial Intelligence and Statistics*, pp. 1082–1092, 2020.
- [14] R. Xu, L. Chen, and A. Karbasi, “Meta learning in the continuous time limit,” *arXiv preprint arXiv:2006.10921*, 2020.
- [15] K. Ji, J. Yang, and Y. Liang, “Multi-step model-agnostic meta-learning: Convergence and improved algorithms,” *arXiv preprint arXiv:2002.07836*, 2020.
- [16] L. Wang, Q. Cai, Z. Yang, and Z. Wang, “On the global optimality of model-agnostic meta-learning,” in *International Conference on Machine Learning*, pp. 9837–9846, PMLR, 2020.
- [17] O. Bousquet and A. Elisseeff, “Stability and generalization,” *Journal of machine learning research*, vol. 2, no. Mar, pp. 499–526, 2002.
- [18] M. Hardt, B. Recht, and Y. Singer, “Train faster, generalize better: Stability of stochastic gradient descent,” in *International Conference on Machine Learning*, pp. 1225–1234, PMLR, 2016.
- [19] A. Rajeswaran, C. Finn, S. M. Kakade, and S. Levine, “Meta-learning with implicit gradients,” in *Advances in Neural Information Processing Systems* (H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, eds.), vol. 32, pp. 113–124, Curran Associates, Inc., 2019.
- [20] L. Collins, A. Mokhtari, and S. Shakkottai, “Task-robust model-agnostic meta-learning,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [21] V. Likhoshervstov, X. Song, K. Choromanski, J. Davis, and A. Weller, “Ufo-blo: Unbiased first-order bilevel optimization,” *arXiv preprint arXiv:2006.03631*, 2020.
- [22] T. Chen, Y. Sun, and W. Yin, “Solving stochastic compositional optimization is nearly as easy as solving stochastic optimization,” *arXiv preprint arXiv:2008.10847*, 2020.
- [23] Y. Hu, S. Zhang, X. Chen, and N. He, “Biased stochastic gradient descent for conditional stochastic optimization,” *arXiv*, vol. abs/2002.10790, 2020.
- [24] C. Finn, A. Rajeswaran, S. Kakade, and S. Levine, “Online meta-learning,” in *Proceedings of the 36th International Conference on Machine Learning*, pp. 97–107, Proceedings of Machine Learning Research (Long Beach, California, USA), pp. 1920–1930, PMLR, 09–15 Jun 2019.
- [25] A. Fallah, A. Mokhtari, and A. Ozdaglar, “Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [26] H. Liu, R. Socher, and C. Xiong, “Taming maml: Efficient unbiased meta-reinforcement learning,” in *International Conference on Machine Learning*, pp. 4061–4071, PMLR, 2019.
- [27] A. Fallah, K. Georgiev, A. Mokhtari, and A. Ozdaglar, “Provably convergent policy gradient methods for model-agnostic meta-reinforcement learning,” *arXiv preprint arXiv:2002.05135*, 2020.
- [28] J. Chen, X.-M. Wu, Y. Li, Q. Li, L.-M. Zhan, and F.-I. Chung, “A closer look at the training strategy for modern meta-learning,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [29] S. Guiroy, V. Verma, and C. Pal, “Towards understanding generalization in gradient-based meta-learning,” *arXiv preprint arXiv:1907.07287*, 2019.
- [30] A. Rakhlin, O. Shamir, and K. Sridharan, “Making gradient descent optimal for strongly convex stochastic optimization,” *arXiv preprint arXiv:1109.5647*, 2011.

- [31] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning* vol. 69, no. 2-3, pp. 169–192, 2007.
- [32] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM Journal on Optimization* vol. 19, no. 4, pp. 1574–1609, 2009.
- [33] R. Bassily, V. Feldman, K. Talwar, and A. Guha Thakurta, "Private stochastic convex optimization with optimal rates," *Advances in Neural Information Processing Systems* vol. 32, pp. 11282–11291, 2019.
- [34] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," *arXiv preprint arXiv:1803.02999*, 2018.
- [35] F. Den Hollander, "Probability theory: The coupling method," lecture notes available online (<http://websites.math.leidenuniv.nl/probability/lecturenotes/CouplingLectures.pdf>), 2012.
- [36] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, vol. 87. Springer, 2004.

Appendix

A Intermediate Results

In this section we list a number of results that will be helpful in proofs of our main results.

Lemma 2 (From [36] with modifications). Let f be a μ -strongly convex and L -smooth function which its gradient is bounded by G over the convex and closed set W . Then, we have

$$\frac{\mu}{2} \|w - w^*\|^2 - f(w) - f(w^*) \leq \frac{L}{2} \|w - w^*\|^2 + G \|w - w^*\| \quad (12)$$

Proof. Recalling the definition of strong convexity and smoothness, we have

$$\frac{\mu}{2} \|w - w^*\|^2 + r(w) \geq f(w) - f(w^*) \leq \frac{L}{2} \|w - w^*\|^2 + G \|w - w^*\| + r(w) \geq f(w) - f(w^*) \quad (13)$$

Since $w^* = \arg \min_W f$, we have $r(w^*) \geq r(w)$, and hence from the left hand side of (13), we immediately obtain the left hand side of (12). To obtain the right hand side, it just suffices to use the bounded gradient assumption along with Cauchy–Schwarz inequality:

$$r(w) \geq r(w^*) - G \|w - w^*\|$$

□

Lemma 3. Suppose the conditions in Assumption 1 are satisfied. Then, with $\mu=1$, and for any batch D and $z \in Z$, we have

$$r(w) - r(\hat{L}(w; D); z) \leq 2G \quad (14)$$

for any $w \in W$. Furthermore, if we take $v \in W$ as well, we have

$$r(w) - r(\hat{L}(w; D); z) - r(v) + r(\hat{L}(v; D); z) \leq 4G \|w - v\| \quad (15)$$

Proof. First, note that

$$r(w) - r(\hat{L}(w; D); z) \leq \|r(w; z) - r(\hat{L}(w; D); z)\| + L \|\hat{L}(w; D) - w\| \leq (1 + L)G \leq 2G \quad (16)$$

where the first inequality follows from smoothness of $r(\cdot; z)$ for any z , and the second inequality is obtained using the bounded gradient assumption. To (15), let us define $\phi(w) = r(w) - r(\hat{L}(w; D); z)$ for any $w \in W$. Note that

$$\phi(w) - \phi(v) = \int_0^1 \langle \nabla \phi(v + s(w - v)), w - v \rangle ds \quad (17)$$

and hence,

$$\begin{aligned} |\phi(w) - \phi(v)| &\leq \int_0^1 \|\nabla \phi(v + s(w - v))\| \|w - v\| ds \\ &\leq \|w - v\| \int_0^1 \|\nabla^2 \hat{L}(v + s(w - v); D) - \nabla^2 \hat{L}(v + s(w - v); D); z\| ds \\ &\leq 2\|w - v\| \int_0^1 \|\nabla^2 \hat{L}(v + s(w - v); D)\| ds \end{aligned} \quad (18)$$

where the last inequality follows from $\|\nabla^2 \hat{L}(v + s(w - v); D)\| \leq L$ and $\mu=1$. Therefore, it suffices to bound

$$\|\nabla^2 \hat{L}(v + s(w - v); D); z\|$$

Using the fact that W is convex, we have $v + s(w - v) \in W$, and hence we could use the same approach in (16) and complete the proof. □

Lemma 4. Suppose Assumptions 1 and 3 hold. Then, with $1=L$, and for any batch D and $z \in \mathbb{Z}$, we have

$$\| \nabla_w \hat{L}(w; D); z \| \leq M + 2G^2; \quad (19)$$

for any $w \in \mathcal{W}$.

Proof. Let $h(\cdot) := \nabla_w \hat{L}(w; D); z$. Using Lemma 3, it is easy to verify that $\|h(\cdot)\| \leq 2G^2$, and hence, using Mean-value Theorem, we have $\|h(0)\| \leq 2G^2$. This result, along with the fact that $\|h(0)\| = \|\nabla_w \hat{L}(w; z)\| \leq M$ by Assumption 3 completes the proof. \square

As we stated in Section 2, MAML uses an unbiased gradient estimate at each iteration. The next lemma provides an upper bound on the variance of such estimate.

Lemma 5. Consider the function $F_i(\cdot; S_i)$ defined in (4) with $\frac{1}{L}$. Suppose the conditions in Assumption 1 are satisfied. Recall that for batch D_i^{in} with size K and D_i^{out} with size b ,

$$g_i(w; D_i^{\text{in}}, D_i^{\text{out}}) = I_d \nabla_w \hat{L}(w; D_i^{\text{in}}) - \nabla_w \hat{L}(w; D_i^{\text{in}}); D_i^{\text{out}}$$

is an unbiased estimate of $F_i(\cdot; S_i)$. Then, for any $w \in \mathcal{W}$, we have

$$\begin{aligned} \|g_i(w; D_i^{\text{in}}, D_i^{\text{out}})\| &\leq 4G; \\ E_{D_i^{\text{in}}, D_i^{\text{out}}} \|g_i(w; D_i^{\text{in}}, D_i^{\text{out}}) - F_i(w; S_i)\|^2 &= O(1)G^2 \left(\frac{2L^2}{K} + \frac{1}{b} \right); \end{aligned}$$

Proof. Recall from Lemma 3 that

$$\| \nabla_w \hat{L}(w; D_i^{\text{in}}); z \| \leq 2G \quad (21)$$

As a result, we have

$$\begin{aligned} \|g_i(w; D_i^{\text{in}}, D_i^{\text{out}})\| &\leq \|I_d \nabla_w \hat{L}(w; D_i^{\text{in}})\| + \|\nabla_w \hat{L}(w; D_i^{\text{in}}); D_i^{\text{out}}\| \\ &\leq (1+L)2G \leq 4G; \end{aligned}$$

To show the second result, we first claim

$$E_{D_i^{\text{in}}} \|g_i(w; D_i^{\text{in}}, S_i^{\text{out}}) - g_i(w; S_i^{\text{in}}, S_i^{\text{out}})\|^2 \leq 36 \frac{2L^2 G^2}{K}. \quad (22)$$

To show this, let us define

$$\begin{aligned} e_H &:= I_d \nabla_w \hat{L}(w; D_i^{\text{in}}) - I_d \nabla_w \hat{L}(w; S_i^{\text{in}}) = \nabla_w \hat{L}(w; S_i^{\text{in}}) - \nabla_w \hat{L}(w; D_i^{\text{in}}) \\ e_G &:= \nabla_w \hat{L}(w; D_i^{\text{in}}); S_i^{\text{out}} - \nabla_w \hat{L}(w; S_i^{\text{in}}); S_i^{\text{out}}; \end{aligned}$$

Note that, by Assumption 1, we have

$$\|e_H\| \leq 2L; \quad \|e_G\| \leq L \|\nabla_w \hat{L}(w; D_i^{\text{in}}) - \nabla_w \hat{L}(w; S_i^{\text{in}})\| \leq 2LG. \quad (24)$$

In addition, using the fact that batch D_i^{in} is chosen uniformly at random, we have

$$E_{D_i^{\text{in}}} [\|e_H\|^2] \leq \frac{2L^2}{K} \frac{n}{n-1}; \quad E_{D_i^{\text{in}}} [\|e_G\|^2] \leq 2L^2 \frac{G^2}{K} \frac{n}{n-1}. \quad (25)$$

Next, note that

$$\begin{aligned} g_i(w; D_i^{\text{in}}, S_i^{\text{out}}) &= g_i(w; S_i^{\text{in}}, S_i^{\text{out}}) \\ &= e_H + \nabla_w \hat{L}(w; S_i^{\text{in}}); S_i^{\text{out}} + e_G + e_H; \end{aligned}$$

Hence, using Cauchy-Schwarz inequality along with (21), we have

$$\begin{aligned} E_{D_i^{in}} \mathbb{E}^h [g_i(w; D_i^{in}; S_i^{out}) - g_i(w; S_i^{in}; S_i^{out})]^2 &\leq \\ &\leq 3(2G)^2 E_{D_i^{in}} [k_{e_H} k^2] + 3(1 + L)^2 E_{D_i^{in}} [k_{e_G} k^2] + 3 E_{D_i^{in}} [k_{e_H} k^2] \\ &\leq 12G^2 E_{D_i^{in}} [k_{e_H} k^2] + 12 E_{D_i^{in}} [k_{e_G} k^2] + 12 L^2 G^2 E_{D_i^{in}} [k_{e_H} k^2]; \end{aligned}$$

where the last inequality is obtained using (24) and 1. Now, using (24), we have

$$\begin{aligned} E_{D_i^{in}} \mathbb{E}^h [g_i(w; D_i^{in}; S_i^{out}) - g_i(w; S_i^{in}; S_i^{out})]^2 &\leq \\ &\leq 12(2 + L^2) \frac{n}{n-1} \frac{K}{1} \frac{L^2 G^2}{K} \leq 36 \frac{L^2 G^2}{K}; \end{aligned}$$

which is the desired claim. Using this result and (21), we imply

$$\begin{aligned} E_{D_i^{in}; D_i^{out}} [g_i(w; D_i^{in}; D_i^{out}) - r \hat{F}_i(w; S_i)]^2 &\leq \\ &\leq E_{D_i^{in}} [g_i(w; D_i^{in}; S_i^{out}) - r \hat{F}_i(w; S_i)]^2 + \frac{4G^2}{b} \\ &\leq 4 \cdot 36 \frac{L^2 G^2}{K} + \frac{G^2}{b} \end{aligned} \quad (26)$$

and the proof is complete. \square

B Proof of Proposition 1

Recall that

$$w^{t+1} = \mathbb{Y}_w w^t + \mathbb{t} g^t;$$

where $g^t := \frac{1}{r} \mathbb{P}_{i \in \mathcal{B}_t} g_i(w^t; D_i^{t, in}; D_i^{t, out})$ is an unbiased estimate of $\hat{F}(w^t)$. Furthermore, by Lemma 5, we know that $\|g^t\| \leq G := 4G$. Also, recall from Lemma 1 that \hat{F} is μ -strongly convex with $\mu = 8$.

Let \mathcal{F}^t be the σ -field generated by the information up to time t (and not including iteration t , such as the randomness \mathcal{B}_t , etc.) It is worth noting that $\mathbb{E}[g^t | \mathcal{F}^t] = r \hat{F}(w^t)$.

First, we claim that similar to the proof of Lemma 1 in [30], we could show

$$\mathbb{E}[k^{t+1} - w^* k^2] \leq (1 - 2\mu) \mathbb{E}[k^t - w^* k^2] + \frac{2}{\mu} G^2; \quad (27)$$

where w^* is the minimizer of $\hat{F}(\cdot; S)$ over W . To see this, and for the sake of completeness, let us recall the steps of the proof. Note that

$$\begin{aligned} \mathbb{E}^h [k^{t+1} - w^* k^2] &\leq \mathbb{E}^h \left[\mathbb{Y}_w w^t + \mathbb{t} g^t - w^* \right]^2 \\ &= \mathbb{E}^h [w^t - w^*]^2 + \mathbb{E}^h [w^t - \mathbb{t} g^t - w^*]^2 \\ &= \mathbb{E}^h [w^t - w^*]^2 + 2 \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t + \mathbb{E}^h [\mathbb{t} g^t - w^*]^2 \\ &= \mathbb{E}^h [w^t - w^*]^2 + 2 \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t + \mathbb{E}^h [\mathbb{t} g^t - w^*]^2; \end{aligned} \quad (28)$$

$$\begin{aligned} &= \mathbb{E}^h [w^t - w^*]^2 + 2 \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t + \mathbb{E}^h [\mathbb{t} g^t - w^*]^2 \\ &= \mathbb{E}^h [w^t - w^*]^2 + 2 \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t + \mathbb{E}^h [\mathbb{t} g^t - w^*]^2; \end{aligned} \quad (29)$$

where (28) follows from non-expansivity of projection and (29) comes from the fact that $\mathbb{E}^h [w^t - w^*] \mathbb{t} g^t = \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t$ and $\mathbb{E}^h [w^t - w^*] \mathbb{t} g^t = \mathbb{E}^h [w^t - w^*] \mathbb{t} g^t$. Now, having (29), and using $\|g^t\| \leq G$ along with the strong convexity

of \hat{F} , we have

$$\begin{aligned} E^h [w^{t+1} - w^2]^i &= E^h [w^t - w^2]^i + 2 \alpha E^h [\hat{F}(w^t) - \hat{F}(w)] + \frac{\alpha^2}{2} E^h [w^t - w^2]^2 + \alpha^2 \mathbb{G}^2 \\ &= E^h [w^t - w^2]^i + 2 \alpha E^h [\hat{F}(w^t) - \hat{F}(w)] + \frac{\alpha^2}{2} E^h [w^t - w^2]^2 + \alpha^2 \mathbb{G}^2 \\ &= (1 - 2\alpha) E^h [w^t - w^2]^i + \alpha^2 \mathbb{G}^2; \end{aligned} \quad (30)$$

where (30) follows from Lemma 2. Next, note that α is given by

$$\alpha = \begin{cases} \frac{1}{t+1}; & \text{for } t \leq \frac{1}{c}; \\ \frac{1}{t}; & \text{for } t > \frac{1}{c}; \end{cases} \quad \text{with } t := b - c;$$

For any t , from (27) and Lemma 2 in [30], we obtain

$$E[kw^t - w^k]^2 \leq \frac{\mathbb{G}^2}{2} + 2\mathbb{G}^2 t \frac{\mathbb{G}^2(t+3)}{2(t+1)}. \quad (31)$$

Also, note that, for $t \leq \frac{1}{c}$, we have

$$E[kw^{t+1} - w^k]^2 \leq (1 - \frac{2}{t+1}) E[kw^t - w^k]^2 + \frac{\mathbb{G}^2}{2(t+1)^2}. \quad (32)$$

Hence, by induction, it can be seen that for any t , we have

$$E[kw^t - w^k]^2 \leq \frac{\mathbb{G}^2(t+3)}{2(t+1)}. \quad (33)$$

Using Lemma 2 gives us (7).

To obtain the bound on the time-average iterate, first, we could similarly, modify the result in [30] to obtain

$$2E[\hat{F}(w^T) - \hat{F}(w)] \leq \frac{1}{T} [kw^0 - w^k]^2 \left(\frac{1}{1} \right) + \sum_{t=1}^{T-1} E[kw^t - w^k]^2 \left(\frac{1}{t+1} - \frac{1}{t} \right) + \mathbb{G}^2 \sum_{t=1}^{T-1} \frac{1}{t^2}.$$

It can be easily verified that for $t = \min(\frac{1}{c}, \frac{8}{t+1})$, the term $\frac{1}{t+1} - \frac{1}{t}$ is always non-positive. Hence, we have

$$\begin{aligned} E[\hat{F}(w^T) - \hat{F}(w)] &\leq [kw^0 - w^k]^2 \frac{1}{T+1} + \frac{2\mathbb{G}^2}{T+1} \sum_{t=0}^{T-1} \frac{1}{t^2} \\ &\leq O(1) \frac{\mathbb{G}^2}{T} (1 + \log(T)) \leq O(1) \frac{\mathbb{G}^2}{T} (1 + \log(T)); \end{aligned} \quad (34)$$

where the last inequality follows from the fact that $\frac{1}{t^2} \leq \frac{1}{t}$ (see Lemma 2 in [30] for the proof.)

C Proof of Theorem 1

To show the claim, it just suffices to show that for any i , we have

$$E_{A;S} [F_i(w_S)] = \hat{F}_i(w_S; S_i). \quad (35)$$

Consider

$$S_1^{\text{in}} = f(z_1^{\text{in}}, \dots, z_n^{\text{in}}; g); \quad S_1^{\text{out}} = f(z_1^{\text{out}}, \dots, z_n^{\text{out}}; g);$$

To see this, first note that

$$F_i(w_S) = E_{f_{z_j, g_{j=1}^K}; z} [w_S] = \hat{L}(w_S; f_{z_j, g_{j=1}^K}; z);$$

where $z_j^{K_{j=1}}$ are K distinct points sampled from \mathcal{D} , and z is also independently sampled from \mathcal{D} . By Assumption 2, we could assume z is different from K other points. Note that we have

$$E_S[F_i(w_S)] = E_{S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i : \quad (36)$$

Next, note that, we can write $F_i(w_S; S_i)$ as

$$F_i(w_S; S_i) = \frac{1}{K |S_i^{out}|} \int_{f_j g_{j=1}^K [n]} \int_{z^{out} [n]} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j^{in} g_{j=1}^K); z^{out} : \quad (37)$$

Thus, we have

$$E_{A;S}[F_i(w_S; S_i)] = \frac{1}{K |S_i^{out}|} \int_{f_j g_{j=1}^K [n]} \int_{z^{out} [n]} \int_{\mathcal{W}_S} E_{A;S} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j^{in} g_{j=1}^K); z^{out} : \quad (38)$$

Notice that, $f_j g_{j=1}^K$ are all different, and hence, due to the symmetry, all the expectations on the RHS are equal. Hence, for a fixed $f_j g_{j=1}^K [n]$ and $z^{out} [n]$, we have

$$\begin{aligned} E_{A;S}[F_i(w_S; S_i)] &= E_{A;S} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j^{in} g_{j=1}^K); z^{out} \\ &= E_{A;S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j^{in} g_{j=1}^K); z^{out} \end{aligned} \quad (37)$$

Next, define the dataset \mathcal{S} by substituting z_j^{in} with z_j , for all j , and z^{out} with z . It is straightforward to see that

$$E_{A;S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j^{in} g_{j=1}^K); z^{out} = E_{A;S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i$$

Therefore, using (37), we obtain

$$E_{A;S}[F_i(w_S; S_i)] = E_{A;S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i : \quad (38)$$

Putting (36) and (38) together, we have

$$\begin{aligned} E_{A;S} F_i(w_S) &= E_{A;S} F_i(w_S; S_i) \\ &= E_{S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i \\ &= E_{S;f z_j g_{j=1}^K; z} \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i \end{aligned} \quad (39)$$

where the last equality follows from Tonelli' theorem. Finally, note that since $(\cdot; K)$ -uniformly stable, we could bound the the inner integral by e.,

$$E_A \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i \leq \int_{\mathcal{W}_S} r \hat{L}(w_S; f z_j g_{j=1}^K); z^i ;$$

and thus, we obtain the desired result (35).

D Proof of Theorem 2

The stability definition says there is ϵ such that the two datasets \mathcal{S} and \mathcal{S}' differ only in the the two following terms:

- S_i^{in} differs from $S_i'^{in}$ in at most K points. We show those K samples by $z_j g_{j=1}^K$ and $z_j' g_{j=1}^K$, respectively.
- S_i^{out} differs from $S_i'^{out}$ in at most one point. We show those by z and z' , respectively.

Let's consider two parallel processes of generating iterates w^t and w^{t*} by using datasets \mathcal{S} and \mathcal{S}^* , respectively. We use the tilde superscript to refer to the second process throughout the proof. Also, we use $D_i^{t;out}$ and $D_i^{t;in}$ to refer to indices of samples $D_i^{t;out}$ and $D_i^{t;in}$, respectively. Also, with a slight abuse of notation, by $\mathbb{E}(w^t; D_i^{in=out})$ we mean $\mathbb{E}(w^t; D_i^{in=out})$.

Note that the randomness of algorithm comes from the randomness in drawing batches at each iteration. We do a coupling argument here. We could assume the two parallel processes of generating iterates w^t and w^{t*} use the same random machine for sampling batches. In other words, $D_i^{t;out} = \tilde{D}_i^{t;out}$, and $D_i^{t;in} = \tilde{D}_i^{t;in}$.

For one particular realization:

- Let u_t be the number of times that the index corresponding to sample (i, z_i) is chosen in $D_i^{t;out}$. Note that this number could be zero in B_t , and it could be greater than one if $i \in B_t$ since $D_i^{t;out}$ is chosen with replacement.
- Let v_t be the number of indices corresponding to the samples $\{(i, z_i)\}_{i=1}^K$ (or $\{z_i\}_{i=1}^K$) that appears in $D_i^{t;in}$. Again, this number could be zero in B_t . Also, note that we take $D_i^{t;in}$ as a batch of K different samples from \mathcal{S}_i^{in} , and hence, each one of indices appears at most one time in $D_i^{t;in}$.

The rest of the proof has three steps:

1. First, recall the definition of u_t and v_t from Algorithm 1. We claim

$$\mathbb{E}[u_t] = \frac{br}{nm}; \quad \mathbb{E}[v_t] = \frac{K^2 r}{nm} \quad (40)$$

The first one is easy to see. $i \in B_t$ with probability r/m , and if that happens, then u_t would have a binomial distribution with mean br/n . To see the second one, note that

$$P(v_t = j) = \binom{K}{j} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j};$$

and therefore,

$$\mathbb{E}[v_t | j \in B_t] = \sum_{j=0}^K \frac{1}{K} \binom{K}{j} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j};$$

Using the fact that $\binom{K}{j} = \frac{K}{j} \binom{K-1}{j-1}$, we obtain

$$\begin{aligned} \mathbb{E}[v_t | j \in B_t] &= \sum_{j=0}^K \frac{K}{K} \binom{K-1}{j-1} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j} \\ &= \sum_{j=0}^K \binom{K-1}{j-1} \frac{1}{j} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j}; \end{aligned} \quad (41)$$

However, note that $\binom{K-1}{j-1} \frac{1}{j} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j}$ is exactly the probability of $v_t = j$ if $K \rightarrow K-1$ and $n \rightarrow n-1$. Hence, the sum $\sum_{j=0}^K \binom{K-1}{j-1} \frac{1}{j} \left(\frac{n}{K}\right)^j \left(\frac{K-j}{K}\right)^{K-j}$ is equal to $\frac{n}{K-1}$, and plugging this into (41) gives us the second part of the claim (40).

2. Second, we claim that under Assumption 1 we have

$$\mathbb{E}_A [kw^T - w^{T*}k] \leq \frac{4G}{mn} (1 + LK) \frac{16(2L + G) + G}{(2L + G)}; \quad (42)$$

Before showing its proof, note that since $\frac{16(2L + G) + G}{(2L + G)} \leq 2$, this could be simplified as

$$\mathbb{E}_A [kw^T - w^{T*}k] \leq O(1) \frac{G}{mn} (1 + LK); \quad (43)$$

Now, let's show why this is true. To simplify the notation, let us define $\eta(w; D; z) := \sum_{i=1}^n w^T r_i \hat{L}(w; D; z)$. We start by revisiting the following lemma from [18]:

Lemma 6. Let ϕ be a μ -strongly convex and L -smooth function. Then, for any $\frac{2}{\mu}$, we have

$$\|k(u - r(u)) - (v - r(v))\| \leq \left(1 + \frac{L}{\mu}\right) \|u - v\|;$$

for any u and v .

Next, recall from Lemma 1 that for any batch \mathcal{B}_t and any $z \in \mathcal{Z}$, $(w; D; z)$ is $4L + 2G$ smooth and $\frac{2}{b}$ strongly convex. Hence, using the above lemma, for \mathcal{B}_t that $\mathcal{B}_t \in \mathcal{B}_t$, we have

$$\|kw_j^{t+1} - w_j^{t+1}\| \leq \left(1 + \frac{2(2L + G)}{16(2L + G) + \mu}\right) \|kw^t - w^t\|; \quad (44)$$

Next, let us assume $\mathcal{B}_t \in \mathcal{B}_t$. In this case, we have

$$\|kw_i^{t+1} - w_i^{t+1}\| \leq \frac{1}{b} \sum_{z \in \mathcal{D}_i^{t, \text{out}}} \|w^t - r(w^t; D_i^{t, \text{in}}; z) - w^t - r(w^t; D_i^{t, \text{in}}; z)\|; \quad (45)$$

$$+ \frac{1}{b} \sum_{z \in \mathcal{D}_i^{t, \text{out}} = \mathcal{D}_i^{t, \text{out}}} \|r(w^t; D_i^{t, \text{in}}; z) - r(w^t; D_i^{t, \text{in}}; z)\|. \quad (46)$$

For (46), note that we know by Lemma 5 that $\|w^t - r(w^t; D; z)\| \leq 4G$, and hence, since $\mathcal{D}_i^{t, \text{out}} = \mathcal{D}_i^{t, \text{out}}$, we could bound the second term by $4G$. As a result, we have

$$\|kw_i^{t+1} - w_i^{t+1}\| \leq \frac{8G}{b} \sum_{z \in \mathcal{D}_i^{t, \text{out}}} \|w^t - r(w^t; D_i^{t, \text{in}}; z) - w^t - r(w^t; D_i^{t, \text{in}}; z)\|; \quad (47)$$

Note that

$$\begin{aligned} & \|w^t - r(w^t; D_i^{t, \text{in}}; z) - w^t - r(w^t; D_i^{t, \text{in}}; z)\| \\ & \|w^t - r(w^t; D_i^{t, \text{in}}; z) - w^t - r(w^t; D_i^{t, \text{in}}; z)\| \\ & + \|r(w^t; D_i^{t, \text{in}}; z) - r(w^t; D_i^{t, \text{in}}; z)\|. \end{aligned} \quad (48)$$

Let us bound the two terms on the RHS (48) separately. First, similar to how we derived 44, we could bound the first term by

$$\begin{aligned} & \|w^t - r(w^t; D_i^{t, \text{in}}; z) - w^t - r(w^t; D_i^{t, \text{in}}; z)\| \\ & \leq \left(1 + \frac{2(2L + G)}{16(2L + G) + \mu}\right) \|kw^t - w^t\|. \end{aligned} \quad (49)$$

To bound the second term on the RHS of (48), note that

$$\begin{aligned} & \|r(w^t; D_i^{t, \text{in}}; z) - r(w^t; D_i^{t, \text{in}}; z)\| \\ & = \left\| \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) - \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) \right\| \\ & \quad + \left\| \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) - \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) \right\| \\ & \quad + \left\| \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) - \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) \right\| \\ & \quad + \left\| \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) - \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) \right\| \\ & \quad + \left\| \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) - \left(\frac{2L}{b} \hat{L}(w^t; D_i^{t, \text{in}}) \right) r(w^t; D_i^{t, \text{in}}; z) \right\| \\ & \leq 2G \|r(w^t; D_i^{t, \text{in}}) - r(w^t; D_i^{t, \text{in}})\|; \end{aligned} \quad (50)$$

where, in the last inequality, we used Lemma 3 along with the third condition of Assumption 1. Hence, what remains is to bound the two terms in (50). To do so, notice that

$$\begin{aligned} & \mathbb{E} \left[\|\mathbf{w}^t - \mathbb{E}[\mathbf{w}^t; \mathcal{D}_i^{t, \text{in}}]; \mathbf{z}\| \right] \leq \mathbb{E} \left[\|\mathbf{w}^t - \mathbb{E}[\mathbf{w}^t; \mathcal{D}_i^{t, \text{in}}]; \mathbf{z}\| \right] \\ & \leq L \mathbb{E} \left[\|\mathbf{w}^t - \mathbb{E}[\mathbf{w}^t; \mathcal{D}_i^{t, \text{in}}]\| \right] \leq 2LG \frac{V_t}{K}; \end{aligned} \quad (51)$$

and

$$\mathbb{E} \left[\|\mathbf{w}^t - \mathbb{E}[\mathbf{w}^t; \mathcal{D}_i^{t, \text{in}}]\|^2 \right] \leq 2L \frac{V_t}{K}. \quad (52)$$

By plugging (51) and (52) into (50) and using $\epsilon = 1$, we have

$$\mathbb{E} \left[\|\mathbf{w}^t - \mathbb{E}[\mathbf{w}^t; \mathcal{D}_i^{t, \text{in}}]; \mathbf{z}\| \right] \leq 8LG \frac{V_t}{K}. \quad (53)$$

Substituting this bound and (49) into (48) and plugging the result into (47), we have

$$\|\mathbf{w}_i^{t+1} - \mathbf{w}_i^{t+1}\| \leq \left(1 - \frac{2(2L + G)}{16(2L + G) + \epsilon}\right) \|\mathbf{w}^t - \mathbf{w}^{t+1}\| + 8G \left(\frac{u_t}{b} + L \frac{V_t}{K}\right). \quad (54)$$

Using (54) and (44), we obtain

$$\frac{1}{r} \sum_{j \in \mathcal{B}_t} \|\mathbf{w}_j^{t+1} - \mathbf{w}_j^{t+1}\| \leq \left(1 - \frac{2(2L + G)}{16(2L + G) + \epsilon}\right) \|\mathbf{w}^t - \mathbf{w}^{t+1}\| + 8G \left(\frac{u_t}{rb} + L \frac{V_t}{rK}\right).$$

Since projections are non-expansive, we have

$$\|\mathbf{w}^{t+1} - \mathbf{w}^{t+1}\| \leq \left(1 - \frac{2(2L + G)}{16(2L + G) + \epsilon}\right) \|\mathbf{w}^t - \mathbf{w}^{t+1}\| + 8G \left(\frac{u_t}{rb} + L \frac{V_t}{rK}\right). \quad (55)$$

Taking an expectation from both sides and using (40), we get

$$\mathbb{E}_A \left[\|\mathbf{w}^{t+1} - \mathbf{w}^{t+1}\| \right] \leq \left(1 - \frac{2(2L + G)}{16(2L + G) + \epsilon}\right) \mathbb{E}_A \left[\|\mathbf{w}^t - \mathbf{w}^{t+1}\| \right] + 8 \frac{\epsilon G}{mn} (1 + LK). \quad (56)$$

Note that we can rewrite this bound as

$$\mathbb{E}_A \left[\|\mathbf{w}^{t+1} - \mathbf{w}^{t+1}\| \right] \leq (1 - \epsilon) \mathbb{E}_A \left[\|\mathbf{w}^t - \mathbf{w}^{t+1}\| \right] + \epsilon;$$

where

$$\epsilon := \frac{2(2L + G)}{16(2L + G) + \epsilon}; \quad \epsilon := \frac{8G}{mn} (1 + LK).$$

Note that the claim (42) is in fact to show

$$\mathbb{E}_A \left[\|\mathbf{w}^t - \mathbf{w}^{t+1}\| \right] \leq \epsilon.$$

This is true for $t = 1$ since $\epsilon \leq \frac{1}{4L+2G} \leq \frac{1}{4}$. Having this, we could easily obtain the result by induction.

3. We are ready to conclude. Note that by Lemma 3, we have

$$\begin{aligned} & \|\mathbf{w}^T - \mathbb{E}[\mathbf{w}^T; f; \mathbf{z}_j; \mathbf{g}_{j=1}^K]; \mathbf{z}\| \leq \|\mathbf{w}^T - \mathbb{E}[\mathbf{w}^T; f; \mathbf{z}_j; \mathbf{g}_{j=1}^K]; \mathbf{z}\| \\ & \leq 4G \|\mathbf{w}^T - \mathbb{E}[\mathbf{w}^T; f; \mathbf{z}_j; \mathbf{g}_{j=1}^K]; \mathbf{z}\| \leq 4(1 + L) \|\mathbf{w}^T - \mathbb{E}[\mathbf{w}^T; f; \mathbf{z}_j; \mathbf{g}_{j=1}^K]; \mathbf{z}\| \\ & \leq 4(1 + L) \|\mathbf{w}^T - \mathbf{w}^T\| \leq 8G \|\mathbf{w}^T - \mathbf{w}^T\| \leq 8G \|\mathbf{w}^T - \mathbf{w}^T\|. \end{aligned}$$

Taking expectations from both sides completes the proof. Note that (43) can be extended to \mathbf{w}^T as well, and using an argument similar to this step, we could show the same stability bound for the average iterates as well.

E Generalization bound for large K regime

Under the premise of Theorem 2, we claim

$$E_{A;S} \left| \frac{1}{n} \sum_{i=1}^n F(w_S) - \hat{F}(w_S; S) \right| \leq O(1) G^2 \left(\frac{1}{mn} + \min \left\{ \frac{LK}{mn}, \frac{1}{K} \right\} \right) \quad (57)$$

To show this, first, recall that

$$F_i(w) = E_{D_i^{\text{test}}} L_i(w) - r \hat{L}(w; D_i^{\text{test}})$$

Let $G_i(w) := L_i(w) - r L_i(w)$: Note that

$$\begin{aligned} |F_i(w) - G_i(w)| &= E_{D_i^{\text{test}}} L_i(w) - r \hat{L}(w; D_i^{\text{test}}) - L_i(w) + r L_i(w) \\ &\leq G E_{D_i^{\text{test}}} \hat{L}(w; D_i^{\text{test}}) - r L_i(w) \leq 4 \frac{G^2}{K} \end{aligned}$$

As a result, for $G(w) = \frac{1}{m} \sum_{i=1}^m G_i(w)$, we have

$$|G(w) - F(w)| \leq O(1) \frac{G^2}{K}$$

Similarly, if we define

$$\hat{G}_i(w) := \hat{L}(w; S_i^{\text{out}}); \quad \hat{G}_w := \frac{1}{m} \sum_{i=1}^m \hat{G}_i(w);$$

we could show that

$$E_S |\hat{G}_w - \hat{F}(w; S)| \leq O(1) \frac{G^2}{K}$$

Finally, note that the well-known generalization results for strongly convex functions by using classic stability definition (Definition 1) implies (see [18] for details)

$$E_{A;S} |G(w_A) - \hat{G}(w_A)| \leq O(1) \frac{G^2}{mn};$$

where w_A is MAML output. Putting these bounds together, we obtain $O(1) \left(\frac{G^2}{mn} + \frac{G^2}{K} \right)$. Taking minimum of this and Theorem 2 proves the aforementioned claim.

Finally, it is worth mentioning that while we are not sure whether our bound is tight for the large regime, this is not necessarily the case that the generalization bound improves as K increases. To see this, consider MAML with only one task, i.e., $m = 1$, and the quadratic loss $\ell(w; z) = (w^T x - y)^2$ with $z = (x; y)$. In addition, and to focus on the generalization error coming from test update, we assume we have access to exact gradients for outer loop, i.e.,

$$\hat{F}(w) = \frac{1}{K} \sum_{f, z_i \in D^{\text{in}}} L(w) - \sum_{i=1}^K \frac{1}{K} r \nabla L(w; z_i)$$

Let $\bar{x} = E[x x^T]$ and $\bar{y} = E[xy]$. Also, we denote the estimation of \bar{x} and \bar{y} over D^{in} by \hat{x} and \hat{y} , respectively.

After some simplifications, it can be shown that

$$r \nabla F(w) = w + 2 \bar{x} w + 2 \bar{y} + O(\epsilon^2);$$

$$r \nabla \hat{F}(w) = w + 2 \hat{x} w + 2 \hat{y} + \hat{\epsilon} + O(\epsilon^2);$$

It can be seen that the difference of the two gradients ($\hat{\epsilon}$) does not decrease as K increases.

F Proof of Theorem 3

First, we show the following lemma:

Lemma 7. For any \mathbf{z} and any $\mathbf{w} \in W$, we have

$$\mathbb{E}_{f_{z_j^{m+1}} \rho_{m+1} g_{j=1}^K} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K); \mathbf{z} \right\rangle \right] - \mathbb{E}_{f_{z_j^i} \rho_i g_{j=1}^K} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right] \leq 4 G^2 k_{\rho_{m+1}} k_{\rho_i} k_{TV} : \quad (58)$$

Proof. Note that since ρ_i are non-atomic, we could assume z_j^i are drawn independently. Same story holds for z_j^{m+1} 's. Now, for any j , let us assume $(z_j^i; z_j^{m+1})$ is drawn from a joint distribution ϕ_j and ρ_{m+1} corresponding to the maximal coupling of these distributions, i.e.,

$$z_j^i \sim \rho_i; z_j^{m+1} \sim \rho_{m+1}; P(z_j^i \in z_j^{m+1}) = k_{\rho_i} k_{\rho_{m+1}} k_{TV} :$$

Hence, with probability $(k_{\rho_i} k_{\rho_{m+1}} k_{TV})^t (1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^{K-t}$, we have $z_j^i \in z_j^{m+1}$ for t choices of j (out of $1; \dots; K$).

In addition, similar to the proof of Lemma 4, we could show that

$$\left| \left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K); \mathbf{z} \right\rangle - \left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right| \leq 2 G k \left| \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K) - \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K) \right|_k :$$

Hence, if $z_j^i \in z_j^{m+1}$ for t choices of j , then we have

$$\left| \left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K); \mathbf{z} \right\rangle - \left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right| \leq 4 G^2 \frac{t}{K} :$$

As a result, we have

$$\begin{aligned} & \mathbb{E}_{f_{z_j^{m+1}} \rho_{m+1} g_{j=1}^K} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K); \mathbf{z} \right\rangle \right] - \mathbb{E}_{f_{z_j^i} \rho_i g_{j=1}^K} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right] \\ & \leq \sum_{t=0}^K \binom{K}{t} (k_{\rho_i} k_{\rho_{m+1}} k_{TV})^t (1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^{K-t} 4 G^2 \frac{t}{K} \\ & = 4 G^2 (k_{\rho_i} k_{\rho_{m+1}} k_{TV}) \sum_{t=0}^K \frac{t}{K} \binom{K}{t} (k_{\rho_i} k_{\rho_{m+1}} k_{TV})^t (1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^{K-t} \\ & = 4 G^2 (k_{\rho_i} k_{\rho_{m+1}} k_{TV}) ; \end{aligned} \quad (59)$$

where the last equality follows from the fact that

$$\sum_{t=0}^K \binom{K}{t} (k_{\rho_i} k_{\rho_{m+1}} k_{TV})^t (1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^{K-t} = \sum_{t=0}^K \binom{K}{t} (k_{\rho_i} k_{\rho_{m+1}} k_{TV})^t (1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^{K-t} = (k_{\rho_i} k_{\rho_{m+1}} k_{TV} + 1 - k_{\rho_i} k_{\rho_{m+1}} k_{TV})^K = 1^K = 1 :$$

□

Let's get back to the proof of Theorem 3. For any $i \in [m+1]$ and any \mathbf{z} , let us define

$$X_i(\mathbf{z}) := \mathbb{E}_{f_{z_j^i} \rho_i g_{j=1}^K} \left[\left\langle \mathbf{w}_S, \mathbf{r} \hat{L}(\mathbf{w}_S; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right] :$$

In other words X_i is the loss over data point \mathbf{z} when the model is updated using the distribution of task i . Next, note that

$$F_{m+1}(\mathbf{w}) - F_i(\mathbf{w}) = \mathbb{E}_{f_{z_j^{m+1}} \rho_{m+1} g_{j=1}^K; \mathbf{z}} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^{m+1}} g_{j=1}^K); \mathbf{z} \right\rangle \right] - \mathbb{E}_{f_{z_j^i} \rho_i g_{j=1}^K; \mathbf{z}} \left[\left\langle \mathbf{w}, \mathbf{r} \hat{L}(\mathbf{w}; f_{z_j^i} g_{j=1}^K); \mathbf{z} \right\rangle \right] : \quad (60)$$

Note that by Lemma 4, the term inside expectation is bounded, and hence, by Fubini's theorem, we can cast this term as

$$E_{z \sim p_{m+1}} [X_{m+1}(z)] = E_{z \sim p_i} [X_i(z)] \quad (61)$$

By Lemma 7, we have $|X_i(z) - X_{m+1}(z)| \leq 4G^2 k_{p_i} - p_{m+1} k_{TV}$. Hence, we have

$$F_{m+1}(w) - F_i(w) = E_{z \sim p_{m+1}} [X_{m+1}(z)] - E_{z \sim p_i} [X_i(z)] = E_{z \sim p_{m+1}} [X_{m+1}(z)] - E_{z \sim p_i} [X_{m+1}(z)] + e_{i,m};$$

where $|e_{i,m}| \leq 4G^2 k_{p_i} - p_{m+1} k_{TV}$. As a result, we have

$$F_{m+1}(w) = \frac{1}{m} \sum_{i=1}^m F_i(w) - E_{z \sim p_{m+1}} [X_{m+1}(z)] + \frac{1}{m} \sum_{i=1}^m E_{z \sim p_i} [X_{m+1}(z)] + 4G^2 k_{p_i} - p_{m+1} k_{TV} \quad (62)$$

Using Lemma 4, we have $|E_{z \sim p_i} [X_{m+1}(z)] - E_{z \sim p_{m+1}} [X_{m+1}(z)]| \leq (M + 2G^2)k_{p_{m+1}} - p_i k_{TV}$. Hence, by (8), we have

$$E_{z \sim p_{m+1}} [X_{m+1}(z)] = \frac{1}{m} \sum_{i=1}^m E_{z \sim p_i} [X_{m+1}(z)] - (M + 2G^2)k_{p_{m+1}} + \frac{1}{m} \sum_{i=1}^m p_i k_{TV} \quad (63)$$

Plugging (63) into (62) gives us the desired result.

F.1 Proof of Corollary 2

Note that

$$E_{A;S} [F_{m+1}(w_S)] = \min_W F_{m+1} - E_{A;S} [F_{m+1}(w_S) - F(w_S)] + E_{A;S} [F(w_S)] = \min_W F + \min_W F_{m+1};$$

where the second term on the right hand side is bounded by assumption, and the first and last term are both bounded by $(p_{m+1}; f_{p_i} g_{j=1}^m)$ based on Theorem 3.

F.2 Generalization to a task drawn from a distribution of recurring and unseen tasks

Here we show how our result for generalization to an unseen task can be extended to the case that the task at test time is generated from a distribution over both recurring tasks $\mathcal{T}_{i=1}^m$ and the unseen task \mathcal{T}_{m+1} .

Corollary 3. Under the premise of Theorem 3, and if the task at the test time is generated from the distribution over $\mathcal{T}_{i=1}^{m+1}$, we have

$$|E [F_i(w)] - F(w)| \leq (T_{m+1}) D(p_{m+1}; f_{p_j} g_{j=1}^m) + (1 - T_{m+1}) \sum_{i=1}^m (T_i) \frac{1}{m} |E [F_i(w)] - F(w)| \leq \frac{1}{m} \sum_{i=1}^m D(p_i; f_{p_j} g_{j=1}^m);$$

where (T_i) is the probability of task \mathcal{T}_i according to distribution.

Proof. Note that

$$|E [F_i(w)] - F(w)| \leq \sum_{j=1}^m |F_{m+1} - F(w)| + (1 - T_{m+1}) \sum_{i=1}^m |E [F_i(w)] - F(w)| \quad (64)$$

Note that by Theorem 3 we have

$$|F_{m+1} - F(w)| \leq D(p_{m+1}; f_{p_j} g_{j=1}^m); \quad |E [F_i(w)] - F(w)| \leq D(p_i; f_{p_j} g_{j=1}^m);$$

Plugging these into (64) completes the proof. \square

G Limitations of the algorithmic stability analysis

Upon reviewers’ suggestion, we briefly discuss why the algorithmic stability technique does not lead to meaningful generalization results for nonconvex loss functions. The main issue with applying the stability framework for the nonconvex case is that we have to select a small stepsize to obtain reasonable generalization bounds, but with such small stepsizes, we cannot guarantee that we will find a first-order stationary point (FOSP) solution of the empirical loss in polynomial time.

To be more precise, consider Theorem 3.12 in Section 3.5 of [18]. There, the authors assume the stepsize α_t satisfies the condition $\alpha_t \leq c/t$. To see how this prohibits us from finding an FOSP efficiently, let us recall the convergence analysis of a non-convex smooth objective function f . There the main inequality is the following (see Section 1.2.3 in [18]):

$$f(w_T) - f^* \geq \sum_{t=0}^{\mathcal{X}} \alpha_t(1 - \alpha_t L/2) \|\nabla f(w_t)\|^2,$$

where L is the smoothness parameter and w_t is t -th iterate. It can be shown that by setting the stepsize to $\alpha_t = \Theta(1/t)$, as suggested by [18], we would require $\exp(\Theta(1/\epsilon^2))$ iterations to find an ϵ -FOSP. However, with a constant stepsize, we can achieve the significantly improved rate of $\mathcal{O}(1/\epsilon^2)$ which matches the lower bound for this setting. As this argument shows, to obtain a meaningful generalization bound using algorithmic stability the stepsize should be selected much smaller than the required threshold and as a result the overall iteration/sample complexity could be very large.

Considering this discussion, the algorithmic stability technique imposes a very restrictive assumption on the stepsizes in the *nonconvex setting* which has a detrimental effect on the training error analysis.

H A toy example

In this section, we provide a simple numerical experiment to validate our theoretical results. We consider a linear regression problem with dimension $d = 10$ for the case that we have m tasks and n samples per task. For each task i , the feature vector x is drawn according to a normal distribution of $\mathcal{N}(\mu_i, 0.2I_d)$, where μ_i is a vector uniformly at random drawn from $[0, 1]^d$. In addition, for a given x , the label y is given by $y = a_i^\top x + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, 0.1)$ and a_i is a random vector. To make tasks similar, we generate the vectors a_i according to $a_i = \frac{u_i + 1_d}{\|u_i + 1_d\|}$, where u_i is a random vector, uniformly drawn from $[0, 1]^d$, and 1_d is the all-one vector.

For the loss function, we consider quadratic loss with quadratic regularization, i.e., $l(w, (x, y)) = (w^\top x - y)^2 + \lambda \|w\|^2$, with $\lambda = 0.01$. We choose the number of samples in the stochastic gradient for adaptation as $K = 5$ and the test time learning rate $\alpha = 0.1$, and run MAML for $T = 20000$ iterations.

Figure 2 shows the dependence of test error over recurring tasks on m and n . In this case the task at test time is a recurring task. We see that the error decreases as m or n increases which is consistent with our theoretical results.

Next, we consider the case that the task at test time is new and unseen. Note that, in this case, from our theoretical results we know that the error bound includes a term $D(p_{m+1}, \{p_i\}_{i=1}^m)$ which does not decay with n . However, if the distributions are close, this term could be relatively small if m is sufficiently large. To study this matter in our example, we consider two cases:

- First, we assume this new task is similar to the observed tasks in training. More formally, similar to the first m tasks, we take $a_{m+1} = \frac{u_{m+1} + 1_d}{\|u_{m+1} + 1_d\|}$, where u_{m+1} is again a random vector, uniformly drawn from $[0, 1]^d$. Figure 3 shows the test error in this case. As we expected, here we do not gain that much from increasing n , but the error decreases as m increases. This matches our intuition, as for small m , i.e., $m = 1$, the distance between two distributions p_1 and p_2 could be large. However, as m increases, we have tasks where their distributions are close to p_{m+1} , and hence the average distance between distributions p_i, \dots, p_m and p_{m+1} decreases.

