Learning-to-learn non-convex piecewise-Lipschitz functions

Anonymous Author(s)
Affiliation
Address
email

Abstract

We analyze the meta-learning of the initialization and step-size of learning algorithms for piecewise-Lipschitz functions, a non-convex setting with applications to both machine learning and algorithms. Starting from recent regret bounds for the exponential forecaster on losses with dispersed discontinuities, we generalize them to be initialization-dependent and then use this result to propose a practical meta-learning procedure that learns both the initialization and the step-size of the algorithm from multiple online learning tasks. Asymptotically, we guarantee that the average regret across tasks of scales with a natural notion of task-similarity that measures the amount of overlap between near-optimal regions of different tasks. Finally, we instantiate the method and its guarantee in two important settings: robust meta-learning and multi-task data-driven algorithm design.

1 Introduction

While learning-to-learn or *meta-learning* has long been an object of study [41], in recent years it has gained significant attention as a multi-task paradigm for developing algorithms for learning in dynamic environments, from multiple sources of data, and in federated settings. Such methods focus on using data gathered from multiple tasks to improve performance when faced with data from a new, potentially related task. Among the more popular approaches to meta-learning is *initialization-based* meta-learning, in which the meta-learner uses multi-task data to output an initialization for an iterative algorithm such as stochastic gradient descent (SGD) [24]. The flexibility of this approach has led to its widespread adoption in areas, such as robotics [22] and federated learning [17], and to a growing number of attempts to understand it, both empirically and theoretically [19, 27, 23, 36, 38]. However, outside some stylized setups our learning-theoretic understanding of how to meta-learn an initialization is largely restricted to the convex Lipschitz setting.

We relax both assumptions to study the meta-learning of online algorithms over piecewise-Lipschitz functions, which can be nonconvex and highly discontinuous. As no-regret online learning over such functions is impossible in-general, we study the case of piecewise-Lipschitz functions whose discontinuities are *dispersed*, i.e. which do not concentrate in any small compact subset of the input domain [10]. Such functions arise frequently in *automated algorithm configuration*, in which the goal is to learn the optimal parameter settings of algorithms for difficult (often NP-Hard) problems over a distribution or sequence of instances; for example, a small change to the metric used to determine cluster linkage can lead to a discontinuous change in the classification error [7]. In this paper, we also demonstrate that such losses are relevant in the setting of adversarial robustness, where we introduce a novel online formulation. For both cases, the associated problems are often solved across many time periods or for many different problem domains, resulting in natural multi-task structure that we might hope to use to improve performance. To the best of our knowledge, ours is the first theoretical study of meta-learning in both of these application settings.

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

In the single-task setting the problem of learning dispersed functions can be solved using simple methods such as the exponentially-weighted forecaster. To design an algorithm for learning to initialize online learners in this setting, we propose a method that optimizes a sequence of data-dependent upper-bounds on the within-task regret [27]. The result is an averaged bound that improves upon the regret of the single-task exponential forecaster so long as there exists an initial distribution that can compactly contain many of the within-task optima of the different tasks. Designing the meta-procedure is especially challenging in our setting because it involves online learning over a set of distributions on the domain. To handle this we study a "prescient" form of the classic follow-the-regularized leader (FTRL) scheme that is run over an unknown discretization; we then show the existence of another algorithm that plays the same actions but uses only known information, thus attaining the same regret while being practical to implement.

To demonstrate the usefulness of our method, we instantiate this algorithm in two settings. In the first we consider learning the right parameter to use for the α-Lloyd's family of clustering algorithms [14]. As the cost function is a piecewise-constant function of α, our meta-procedure can help learn the right distribution over α to initialize the exponential forecaster for this task. In experimental evaluations on two datasets—a synthetic Gaussian mixture model and the well-known Omniglot dataset from meta-learning [30]—our meta-procedure leads to improved clustering accuracy compared to single-task learning to cluster. This result holds for both one-shot and five-shot clustering tasks.

The second instantiation of our meta-learning procedure is to a new notion of adversarial robustness for the setting of online learning, where our results imply robust meta-learning in the presence of outliers. In this setting, the adversary can make (typically small) modifications to some example $x \in \mathcal{X}$, which can result in potentially large changes to the corresponding loss value $l_h(x)$, where $h \in \mathcal{H}$ is our hypothesis. For instance, consider the well-studied setting of adversarial examples for classification of images using deep neural networks [32, 16]. Given a neural network f, the adversary can perturb a datapoint x to a point x', say within a small L_p-ball around x, such that $f(x) = f(x')$ but the true label of x' does not match x, and therefore $l_f(x) \neq l_f(x')$. In general, under the adversarial influence, we observe a *perturbed loss* function $\tilde{l}_h(x) = l_h(x) + a_h(x)$. Typically we are interested in optimizing both the perturbed loss $\tilde{l}_h(x)$, i.e. measuring performance relative to optimum for adversarially perturbed losses, and the *true loss* $l_h(x)$ (performance on the unobserved, unperturbed loss). For example, in the online learning setting, [1] consider perturbed loss minimization for linear dynamical systems, while [37] look at true $\{0, 1\}$ loss minimization in the presence of adversarial noise. Our approach ensures that regret for both the perturbed and true loss are small, for piecewise-Lipschitz but dispersed adversaries.

1.1 Related work

The success of meta-learning has led to significant theoretical effort to understand it. Most efforts studying initialized-based meta-learning focus on the convex Lipschitz setting [19, 28]; work studying inherently nonconvex modeling approaches instead usually study multi-task representation learning [6, 31, 21, 43] or target optimization, e.g. stationary point convergence [23]. An exception is a study of linear models over Gaussian data showing that nonconvexity is critical to meta-learning an initialization that exploits low-rank task structure [38]. There is also work extending results from the neural tangent kernel literature to meta-learning [44], but in this case the objective becomes convex. On the other hand, we study initializations for learning a class of functions that can be highly non-convex and have numerous discontinuities. Theoretically, our work uses the Average Regret-Upper-Bound Analysis (ARUBA) strategy [27] for obtaining a meta-update procedure for initializing within-task algorithms; the main technical advance in our work is in providing the guarantees for it in our setting, which is challenging due to the need to learn over a space of probability measures.

Data-driven configuration is the selection of an algorithm from a parameterized family, by doing learning over multiple problem instances [25, 11]. In other words, it is 'hyperparameter tuning' with formal guarantees, and has applications to integer programming, clustering and learning with limited labeled data [8, 7, 13]. In this work, we show how this general approach can be made even more effective by enabling it to adapt to task similarity. We also show applications of our results to robust meta-learning in the presence of outliers in the dataset [35, 29]. While previous work on robust online learning has considered adversaries with bounded perturbation in the online learning setting [1, 37], our results allow potentially unbounded perturbations provided the adversary uses a smooth distribution. That is, the adversarial attack can be thought of as a distribution of perturbations,

similar to the smoothed analysis approach of [40]. In the offline setting, a similar attack is studied in the context of deep network feature-space attacks by [5]. We also remark that our formulation has a poisoning aspect, since we do not observe the clean loss $l_h(x)$, which is of particular interest in federated learning [3, 42]. Also, note that unlike the typical applications of data-driven design where optimization is over the dual loss function, i.e. loss as a function of the algorithm parameter for a fixed sample $x \in \mathcal{X}$, here we consider learning loss or confidence functions over the input space $\mathcal{X}$.

2 Preliminaries and initialization-dependent learning of dispersed functions

In this section we introduce our setup and notation for online learning of piecewise-Lipschitz functions in a multi-task environment. We then generalize existing results for the single-task setting in order to obtain within-task regret bounds that depend on both the initialization and the task data. This is critical for both defining a notion of task similarity and devising a meta-learning procedure.

2.1 Meta-learning setup

Following past setups [2, 20, 27], for some $T, m > 0$ and all $t \in [T]$ and $i \in [m]$ we consider a meta-learner faced with a sequence of Tm loss functions $\ell_{t,i} : C \mapsto [0, 1]$ over a compact subset $C \subset \mathbb{R}^d$ that lies within a ball $\mathcal{B}(\rho, R)$ of radius R around some point $\rho \in \mathbb{R}^d$. Here we used the notation $[n] = \{1, \ldots, n\}$. Before each loss function $\ell_{t,i}$ the meta-learner must pick an element $\rho_{t,i} \in C$ before then suffering a loss or cost $\ell_{t,i}(\rho_{t,i})$. For a fixed t, the subsequence $\ell_{t,1}, \ldots, \ell_{t,m}$ defines a **task** for which we expect a single element $\rho_t^* \in C$ to do well, and thus we will use the **within-task regret** on task t to describe the quantity

$$\mathbf{R}_{t,m} = \sum_{i=1}^{m} \ell_{t,i}(\rho_{t,i}) - \ell_{t,i}(\rho_t^*) \quad \text{where} \quad \rho_t^* \in \arg\min_{\rho \in C} \sum_{i=1}^{m} \ell_{t,i}(\rho) \tag{1}$$

In the single-task setting the goal is usually to show that $R_{t,m}$ is sublinear in m, i.e. that the average loss decreases with more rounds. A key point here is that the functions we consider can have numerous global optima. In this work we will assume, after going through the m rounds of task t, that we have oracle access to a single fixed optimum for t, which we will refer to using ρ_t^* and use in both our algorithm and to define the task-similarity. Note that in the types of applications we are interested in—piecewise-Lipschitz functions—the complexity of computing optima scales with the number of discontinuities. In the important special case of piecewise-constant functions, this dependency becomes logarithmic [18]. Thus this assumption does not affect the usefulness of the result.

Our goal will be to improve the guarantees for regret in the single-task case by using information obtained from solving multiple tasks. In particular, we expect average performance across tasks to improve as we see more tasks; to phrase this mathematically we define the **task-averaged regret**

$$\bar{\mathbf{R}}_{T,m} = \frac{1}{T} \sum_{t=1}^{T} \mathbf{R}_{t,m} = \frac{1}{T} \sum_{t=1}^{T} \sum_{i=1}^{m} \ell_{t,i}(\rho_{t,i}) - \ell_{t,i}(\rho_t^*) \tag{2}$$

and claim improvement over single-task learning if in the limit of $T \to \infty$ it is smaller than $\mathbf{R}_{t,m}$. Note that for simplicity in this work we assume all tasks have the same number of rounds within-task, but as with past work our results are straightforward to extend to the more general setting.

2.2 Learning piecewise-Lipschitz functions

We now turn to our target functions and within-task algorithms for learning them: piecewise-Lipschitz losses, i.e. functions that are L-Lipschitz w.r.t. the Euclidean norm everywhere except on measure zero subsets of the space; here they may have arbitrary jump discontinuities so long they still bounded between $[0, 1]$. Apart from being a natural setting of interest due to its generality compared to past work on meta-learning, this class of functions has also been shown to have important applications in data-driven algorithm configuration [10]; there these functions represent the cost, e.g. an objective value or time-complexity, of algorithms for difficult problems such as integer programming, auction design, and clustering.

This literature has also shown lower bounds demonstrating that no-regret learning piecewise-Lipschitz function is impossible in general, necessitating assumptions about the sequence. One such condition is *dispersion*, which requires that the discontinuities are not too concentrated.

Algorithm 1 Exponential Forecaster

Input: step size parameter $\lambda \in (0, 1]$, initialization $w : C \to \mathbb{R}_{\geq 0}$.
Initialize $w_1 = w$
for $i = 1, 2, \ldots, m$ **do**
  $W_i := \int_C w_i(\rho) d\rho$
  Sample ρ_i with probability proportional to $w_i(\rho_i)$, i.e. with probability $p_i(\rho_i) = \frac{w_i(\rho_i)}{W_i}$
  Suffer $\ell_i(\rho_i)$ and observe $\ell_i(\cdot)$
  For each $\rho \in C$, set $w_{i+1}(\rho) = e^{-\lambda \ell_i(\rho)} w_i(\rho)$

Definition 2.1 ([10])**.** *The sequence of random loss functions $\ell_1, \ldots, \ell_m$ is said to be β-***dispersed** *with Lipschitz constant L if, for all m and for all $\epsilon \geq m^{-\beta}$, we have that, in expectation over the randomness of the functions, at most $\tilde{O}(\epsilon m)$ functions (the soft-O notation suppresses dependence on quantities beside ϵ, m and β, as well as logarithmic terms) are not L-Lipschitz for any pair of points at distance ϵ in the domain $\mathbb{C}$. That is, for all m and for all $\varepsilon \geq m^{-\beta}$,*

$$\mathbb{E}\left[\max_{\substack{\rho,\rho' \in \mathbb{C} \\ \|\rho-\rho'\|_2 \leq \epsilon}} \left\|\left\{i \in [m] \mid \ell_i(\rho) - \ell_i(\rho') > L\|\rho - \rho'\|_2\right\}\right|\right] \leq \tilde{O}(\epsilon m) \tag{3}$$

Assuming a sequence of m β-dispersed loss functions and initial distribution w_1 set to the uniform distribution over C and optimize the step size parameter, the exponential forecaster presented in Algorithm 1 achieves sublinear regret $\tilde{O}(\sqrt{dm \log(Rm)} + (L+1)m^{1-\beta})$. While this result achieves a no-regret procedure, its lack of dependence on both the task-data and on the chosen initialization makes it difficult to meta-learn. In the following theorem, we generalize the regret bound for the exponential forecaster to make it data-dependent and hyperparameter-dependent:

Theorem 2.1. *Let $\ell_1, \ldots, \ell_m : C \mapsto [0, 1]$ be any sequence of piecewise L-Lipschitz functions that are β-dispersed. Suppose $C \subset \mathbb{R}^d$ is contained in a ball of radius R. The exponentially weighted forecaster (Algorithm 1) has expected regret $\mathbf{R}_m \leq m\lambda + \frac{\log(1/Z)}{\lambda} + \tilde{O}((L+1)m^{1-\beta})$, where $Z = \frac{\int_{\mathcal{B}(\rho^*, m^{-\beta})} w(\rho) d\rho}{\int_C w(\rho) d\rho}$ for ρ^* the optimal action in hindsight.*

The proof of this result adapts past analyses of Algorithm 1; setting step-size λ appropriately recovers the previously mentioned bound. The new bound is useful due to its explicit dependence on both the initialization w and the optimum in hindsight via the $\log(1/Z)$ term. Assuming w is a (normalized) distribution, this effectively measures the overlap between the chosen initialization and a small ball around the optimum; we thus call $-\log Z = -\log \frac{\int_{\mathcal{B}(\rho^*, m^{-\beta})} w(\rho) d\rho}{\int_C w(\rho) d\rho}$ the **negative log-overlap**.

We also obtain an asymptotic lower bound on the expected regret of any algorithm by extending the argument of [9] to the multi-task setting. We show that for finite D^* we must suffer $\tilde{\Omega}(m^{1-\beta})$ regret, which limits the improvement we can hope to achieve from task-similarity.

Theorem 2.2. *There is a sequence of piecewise L-Lipschitz β-dispersed functions $\ell_{i,j} : [0, 1] \mapsto [0, 1]$, whose optimal actions in hindsight $\arg\min_\rho \sum_{i=1}^m l_{t,i}(\rho)$ are contained in some fixed ball of diameter D^*, for which any algorithm has expected regret $\mathbf{R}_m \geq \tilde{\Omega}(m^{1-\beta})$.*

2.3 Task-similarity

Before proceeding to our discussion of meta-learning, we first discuss what we might hope to achieve with it; specifically, we consider what a reasonable notion of task-similarity is in this setting. Note that the Theorem 2.1 regret bound has three terms, of which two depend on the hyperparameters and the last is due to dispersion and cannot be improved via better settings. Our focus will thus be on improving the first two terms, which are the dominant ones due to the dependence on the dimensionality and the distance from the initialization encoded in the negative log overlap. In particular, when the initialization is the uniform distribution then this quantity depends inversely on the size of a small ball around the optimum, which may be quite small. Via meta-learning we hope to assign more of the probability mass of the initializer to areas close to the optimum, which will

decrease these terms. On average, rather than a dependence on the volume of a small ball we aim to achieve a dependence on the **average negative log-overlap**

$$V^2 = -\min_{w:C\mapsto\mathbb{R}_{\geq 0},\int_C w(\rho)d\rho=1} \frac{1}{T}\sum_{t=1}^{T}\log\int_{\mathcal{B}(\rho_t^*,m^{-\beta})} w(\rho)d\rho \tag{4}$$

which can be much smaller if the task optima ρ_t^* are close together; for example, if they are the same then $V = 0$. The settings of Algorithm 1 that achieve this are directly related to this quantity: the optimal initializer is the distribution achieving V and the optimal step-size is $V/\sqrt{m}$. Note that while the explicit definition requires computing a minimum over a set of functions, the task-similarity can be computed using the discretization constructed in Section 3.1.

3 An algorithm for meta-learning the initialization and step-size

Having established a single-task algorithm and shown how its regret depends on the initialization and step-size, we move on to meta-learning these hyperparameters. Recall that our goal is to make the task-averaged regret (2) small, in particular to improve upon the baseline of repeatedly running Algorithm 1 from the uniform distribution, up to $o_T(1)$ terms that vanish as we see more tasks. This accomplishes the meta-learning goal of using multiple tasks to improve upon single-task learning.

In this paper, we use the strategy of running online learning algorithms on the data-dependent regret guarantees from above [27]. If we can do so with sublinear regret in T, then we will improve upon the single-task guarantees up to $o_T(1)$ terms, as desired. Specifically, we are faced with a sequence of regret-upper-bounds $U_t(w,v) = (v + f_t(w)/v)\sqrt{m} + g(m)$ on nonnegative functions w over C and positive scalars $v > 0$. Note that $g(m)$ cannot be improved via meta-learning, so we will focus on learning w and v. To do so, we run two online algorithms, one over the functions f_t and the other over $h_t(v) = v + f_t(w_t)/v$, where w_t is set by the first procedure. As shown in the following lemma, if both procedures have sublinear regret then our task-averaged regret will have the desired properties:

Lemma 3.1. *Assume each task $t \in [T]$ consists of a sequence of m β-dispersed piecewise L-Lipschitz functions $\ell_{t,i} : C \mapsto [0,1]$. Let f_t and g be functions such that the regret of Algorithm 1 run with step-size $\lambda = v\sqrt{m}$ for $v > 0$ and initialization $w : C \mapsto \mathbb{R}_{\geq 0}$ is bounded by $U_t(w,v) = (v + f_t(w)/v)\sqrt{m} + g(m)$. Suppose we have a procedure that achieves $F_T(w)$ regret w.r.t. any $w : C \mapsto \mathbb{R}_{\geq 0}$ by playing actions $w_t : C \mapsto \mathbb{R}_{\geq 0}$ on f_t and another procedure that achieves $H_T(v)$ regret w.r.t. any $v > 0$ by playing actions $v_t > 0$ on $h_t(v) = v + f_t(w_t)/v$, where H_T is non-increasing on the positive reals. Then by setting $\rho_{t,i}$ using Algorithm 1 with step-size $v_t/\sqrt{m}$ and initialization w_t at each task t we get task-averaged regret bounded by*

$$\left(\frac{H_T(V)}{T} + \min\left\{\frac{F_T(w^*)}{VT}, 2\sqrt{F_T(w^*)/T}\right\} + 2V\right)\sqrt{m} + g(m) \tag{5}$$

for $w^* = \arg\min_{w:C\mapsto\mathbb{R}_{\geq 0}} \sum_{t=1}^{T} f_t(w)$ the optimal initialization and V the task-similarity (4).

This result is an analog of [27, Theorem 3.1] and follows by manipulating the definition of regret. It reduces the problem of obtaining a small task-averaged regret to solving two online learning problems, one to set the initialization and one to set the step-size. So long as both have sublinear regret then we will improve over single-task learning. In the next two sections we derive suitable procedures.

3.1 Meta-learning the initialization

We now come to the most technically challenging component of our meta-learning procedure: learning the initialization. As discussed above, we can accomplish this by obtaining a no regret procedure for the function sequence $f_t(w) = -\log\frac{\int_{\mathcal{B}(\rho_t^*,m^{-\beta})} w(\rho)d\rho}{\int_C w(\rho)d\rho}$. This is nontrivial as the optimization domain is a set of nonnegative functions, effectively measures on the domain C. To handle this, we first introduce some convenient notation and abstractions. At each task t we are faced with some function f_t associated with an unknown closed subset $C_t \subset C$—in particular $C_t = \mathcal{B}(\rho_t^*, m^{-\beta})$—with positive volume $\mathrm{vol}(C_t) > 0$ that is revealed after choosing $w_t : C \mapsto \mathbb{R}_{\geq 0}$ For each time t define the discretization $\mathcal{D}_t = \{D = \bigcap_{s\leq t} C_s^{(\mathbf{c}_{[s]})} : \mathbf{c} \in \{0,1\}^t, \mathrm{vol}(D) > 0\}$ of C, where $C_t^{(0)} = C_t$ and $C_t^{(1)} = C\backslash C_t$. We will use elements of these discretizations to index nonnegative vectors in

Algorithm 2 Follow-the-Regularized-Leader (prescient form)

Input: discretization $\mathcal{D}_T$ of C, mixture parameter $\gamma \in [0,1]$, step-size $\eta > 0$
Initialize $\mathbf{w}_1 = \hat{\mathbf{v}}$
for $t = 1, 2, \ldots, T$ **do**
  Play $\mathbf{w}_t$.
  Suffer $f_t(\mathbf{w}_t) = -\log\langle \mathbf{w}_t^*, \mathbf{w}_t\rangle$.
  Observe f_t.
  Update $\mathbf{w}_{t+1} = \arg\min_{\|\mathbf{w}\|_1=1, \mathbf{w}\geq\gamma\hat{\mathbf{v}}} D_{KL}(\mathbf{w}||\hat{\mathbf{v}}) + \eta\sum_{s\leq t} f_s(\mathbf{w})$

$\mathbb{R}_{\geq 0}^{|\mathcal{D}_t|}$; specifically, for any measure $w : C \mapsto \mathbb{R}_{\geq 0}$ let $\mathbf{w}(t) \in \mathbb{R}_{\geq 0}^{|\mathcal{D}_t|}$ denote the vector with entries $\mathbf{w}(t)_{[D]} = \int_D w(\rho)d\rho$ for $D \in \mathcal{D}_t$. Note that we will exclusively use p, q, v, w for measures, with v specifically referring to the uniform measure, i.e. $\mathbf{v}(t)_{[D]} = \mathrm{vol}(D)$. For convenience, for all real vectors $\mathbf{x}$ we will use $\hat{\mathbf{x}}$ to denote $\mathbf{p}/\|\mathbf{p}\|_1$. Finally, we abuse notation and remove the parentheses to refer those vectors associated with the final discretization, i.e. $\mathbf{v} = \mathbf{v}(T)$ and $\mathbf{w} = \mathbf{w}(T)$.

Now that we have this notation we can turn back to the functions we are interested in: $f_t(w) = -\log\frac{\int_{C_t} w(\rho)d\rho}{\int_C w(\rho)d\rho}$, where $C_t = \mathcal{B}(\rho_t^*, m^{-\beta})$. Observe that we can equivalently write this as $f_t(\mathbf{w}) = -\log\langle \mathbf{w}_t^*, \hat{\mathbf{w}}\rangle$, where $\mathbf{w}_{t[D]}^* = 1_{D\subset C_t}$; this translates our online learning problem from the domain of measures on C to the simplex on $|\mathcal{D}_T|$ elements. However, we cannot play in this domain explicitly as we do not have access to the final discretization $\mathcal{D}_T$, nor do we get access to $\mathbf{w}_t^*$ after task t, except implicitly via C_t. In this section we design a method that implicitly run an online convex optimization procedure over $\mathbb{R}_{\geq 0}^{|\mathcal{D}_T|}$ while explicitly playing probability measures $w : C \mapsto \mathbb{R}_{\geq 0}$.

As the functions f_t are exp-concave, one might first consider applying a method attaining logarithmic regret on such losses [26, 33]; however, such algorithms have regret that depends linearly on the dimension, which in our case is $\mathrm{poly}(T)$. We thus turn to the the follow-the-regularized-leader (FTRL) family of algorithms, which in the case of entropic regularization are well-known to have regret logarithmic in the dimension [39]. In Algorithm 2 we display the pseudo-code of a modification with regularizer $D_{KL}(\cdot||\hat{\mathbf{v}})$, where recall $\mathbf{v}$ is the vector of volumes of the discretization $\mathcal{D}_T$ of C, and we constrain the played distribution to have measure at least $\gamma\hat{\mathbf{v}}_{[D]}$ over every set $D \in \mathcal{D}_T$.

While Algorithm 2 explicitly requires knowing the discretization $\mathcal{D}_T$ of C in advance, the following key lemma shows that we can run the procedure knowing only the discretization $\mathcal{D}_t$ after task t by simply minimizing the same objective over probability distributions discretized on $\mathcal{D}_t$. This crucially depends on the re-scaling of the entropic regularizer by $\hat{\mathbf{v}}$ (which notably corresponds to the uniform distribution over C) and the fact that $\mathbf{w}_t^* \in \{0,1\}^{|\mathcal{D}_T|}$.

Lemma 3.2. *Let $w : C \mapsto \mathbb{R}_{\geq 0}$ be the probability measure corresponding to the minimizer*

$$\mathbf{w} = \arg\min_{\|\mathbf{q}\|_1=1, \mathbf{q}\geq\gamma\hat{\mathbf{v}}} D_{KL}(\mathbf{q}||\hat{\mathbf{v}}) - \eta\sum_{s\leq t}\log\langle \mathbf{w}_s^*, \mathbf{q}\rangle \tag{6}$$

and let $\tilde{w} : C \mapsto \mathbb{R}_{\geq 0}$ be the probability measure corresponding to the minimizer

$$\tilde{\mathbf{w}}(t) = \arg\min_{\|\mathbf{q}\|_1=1, \mathbf{q}\geq\gamma\hat{\mathbf{v}}(t)} D_{KL}(\mathbf{q}||\hat{\mathbf{v}}(\mathbf{t})) - \eta\sum_{s\leq t}\log\langle \mathbf{w}_s^*(t), \mathbf{q}\rangle \tag{7}$$

Then $\mathbf{w} = \tilde{\mathbf{w}}$.

We can thus move on to proving a regret guarantee for Algorithm 2. This follows from Jensen's inequality together with standard results for FTRL once we show that the loss functions are $\frac{1}{\gamma\,\mathrm{vol}(C_t)}$-Lipschitz over the constrained domain, yielding the following guarantee for Algorithm 2:

Theorem 3.1. *Algorithm 2 has regret bounded by*

$$\frac{1-\gamma}{\eta}D_{KL}(\mathbf{w}^*||\hat{\mathbf{v}}) + \frac{\eta}{\gamma^2}\sum_{t=1}^{T}\frac{1}{(\mathrm{vol}(C_t))^2} + \gamma\sum_{t=1}^{T}\log\frac{1}{\mathrm{vol}(C_t)} \tag{8}$$

w.r.t. the optimum in hindsight $\mathbf{w}^* \in \arg\min_{\|\mathbf{w}\|_1=1, \mathbf{w}\geq\mathbf{0}}\sum_{t=1}^{T} f_t(\mathbf{w})$ of the functions f_t. Setting $\gamma^2 = GB/\sqrt{T}$ and $\eta^2 = \frac{B^2\gamma^2}{TG^2}$, where $B^2 = D_{KL}(\mathbf{w}^*||\hat{\mathbf{v}})$ and $G^2 = \frac{1}{T}\sum_{t=1}^{T}\frac{1}{(\mathrm{vol}(C_t))^2}$, yields sublinear regret $\tilde{O}(\sqrt{BG}T^{\frac{3}{4}})$.

Proof. Algorithm 2 is standard FTRL with regularizer $\frac{1}{\eta}D_{KL}(\cdot||\hat{\mathbf{v}})$, which has the same Hessian as the standard entropic regularizer over the simplex and is thus $\frac{1}{\eta}$-strongly-convex w.r.t. $\|\cdot\|_1$ [39, Example 2.5]. Applying Jensen's inequality, the standard regret bound for FTRL [39, Theorem 2.11] together with the Lipschitz guarantee of Claim **??**, and Jensen's inequality again yields the result:

$$\begin{aligned}
\sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\mathbf{w}^*) &= \sum_{t=1}^{T} f_t(\mathbf{w}_t) - (1-\gamma)f_t(\mathbf{w}^*) - \gamma f_t(\hat{\mathbf{v}}) + \gamma(f_t(\hat{\mathbf{v}}) - f_t(\mathbf{w}^*)) \\
&\leq \sum_{t=1}^{T} f_t(\mathbf{w}_t) - f_t(\gamma\hat{\mathbf{v}} + (1-\gamma)\mathbf{w}^*) + \gamma \log \frac{\langle \mathbf{w}_t^*, \mathbf{w}^* \rangle}{\langle \mathbf{w}_t^*, \hat{\mathbf{v}} \rangle} \\
&\leq \frac{1}{\eta} D_{KL}(\gamma\hat{\mathbf{v}} + (1-\gamma)\mathbf{w}^* || \hat{\mathbf{v}}) + \frac{\eta}{\gamma^2} \sum_{t=1}^{T} \frac{1}{(\mathrm{vol}(C_t))^2} + \gamma \sum_{t=1}^{T} \log \frac{1}{\mathrm{vol}(C_t)} \\
&\leq \frac{1-\gamma}{\eta} D_{KL}(\mathbf{w}^* || \hat{\mathbf{v}}) + \frac{\eta}{\gamma^2} \sum_{t=1}^{T} \frac{1}{(\mathrm{vol}(C_t))^2} + \gamma \sum_{t=1}^{T} \log \frac{1}{\mathrm{vol}(C_t)}
\end{aligned}$$

□

Since the regret is sublinear in T, this result satisfies our requirement for attaining asymptotic improvement over single-task learning via Lemma 3.1. However, there are several aspects of this bound that warrant some discussion. The first is the rate of $T^{\frac{3}{4}}$, which is less sublinear than the standard $\sqrt{T}$ and certainly the $\log T$ regret of exp-concave functions. However, the functions we face are (a) non-Lipschitz and (b) over a domain that has dimensionality $\Omega(T)$; both violate conditions for good rates in online convex optimization [26, 39], making our problem much more difficult.

A more salient aspect is the dependence on $B^2 = D_{KL}(\mathbf{w}^*||\hat{\mathbf{v}})$, effectively the negative entropy of the optimal initialization. This quantity is in-principle unbounded but is analogous to standard online convex optimization bounds that depend on the norm of the optimum, which in e.g. the Euclidean case are also unbounded. In our case, if the optimal distribution is highly concentrated on a very small subset of the space it will be difficult to compete with. Note that our setting of η depends on knowing or guessing B; this is also standard but is certainly a target for future work to address. For example, past work on parameter-free algorithms has solutions for optimization over the simplex [34]; however, it is unclear whether this is straightforward to do while preserving the property given by Lemma 3.2 allowing us to implicitly work with an unknown discretization. A more reasonable approach may be to compete only with smooth measures that only assign probability at most $\kappa\,\mathrm{vol}(D)$ to any subset $D \subset C$ for some constant $\kappa \geq 1$; in this case we will simply have B bounded by $\log \kappa$.

A final issue is the dependence on $\sqrt{G}$, which is bounded by the reciprocal of the smallest volume $\mathrm{vol}(C_t)$, which in the dispersed case is roughly $O(m^{\beta d})$; this means that the task-averaged regret will have a term that, while decreasing as we see additional tasks, is *increasing* in the number of within-task iterations and the dispersion parameter, which is counter-intuitive. It is also does so exponentially in the dimension. Note that in the common algorithm configuration setting of $\beta = 1/2$ and $d = 1$ this will simply mean that for each task we suffer an extra $o_T(1)$ loss at each within-task round, a quantity which vanishes asymptotically.

3.2 Meta-learning the stepsize

In addition to learning the initialization, Lemma 3.1 requires learning the task-similarity to set the within-task step-size $\lambda > 0$. This involves optimizing functions of form $h_t(v) = v + f_t(w_t)/v$. Since we know that the measures w_t are lower-bounded in terms of γ, we can apply a previous result [27] that solves this by running the EWOO algorithm [26] on the modified sequence $v + \frac{f_t(w_t)+\varepsilon^2}{v}$:

Corollary 3.1. *For any $\varepsilon > 0$, running the EWOO algorithm on the modified sequence $v + \frac{f_t(w)+\varepsilon^2}{v}$ over the domain $[\varepsilon, \sqrt{D^2 - \log\gamma + \varepsilon^2}]$, where $D^2 \geq \frac{1}{T}\sum_{t=1}^{T} \log \frac{1}{\mathrm{vol}(C_t)}$, attains regret*

$$\min\left\{\frac{\varepsilon^2}{v^*}, \varepsilon\right\} T + \frac{\sqrt{D^2 - \log\gamma}}{2} \max\left\{\frac{D^2 - \log\gamma}{\varepsilon^2}, 1\right\} (1 + \log(T+1)) \tag{9}$$

on the original sequence $h_t(v) = v + f_t(w)/v$ for all $v^* > 0$.

Algorithm 3 Meta-learning the parameters of the exponential forecaster (Algorithm 1)

Input: domain $C \subset \mathbb{R}^d$, step-size $\eta > 0$, FTRL param. $\gamma \in [0, 1]$, EWOO param. $\varepsilon, D > 0$.
Initialize w_1 to the uniform measure on C.
Initialize EWOO with exp-concavity $\frac{2}{D} \min\left\{\frac{\varepsilon^2}{D^2}, 1\right\}$ on domain $[\varepsilon, \sqrt{D^2 + \varepsilon^2 - \log \gamma}]$. This sets the first step-size λ_1.
for task $t = 1, 2, \ldots, T$ **do**
  Run Algorithm 1 with initialization w_t and step-size λ_t.
  Set w_{t+1} to $\mathbf{w}_{t+1}(t) = \arg\min_{\|\mathbf{w}\|_1 = 1, \mathbf{w} \geq \gamma \hat{\mathbf{v}}(t)} D_{KL}(\mathbf{w} || \hat{\mathbf{v}}(t)) - \eta \sum_{s \leq t} \log \langle \mathbf{w}_s^*(t), \mathbf{w} \rangle$.
  Pass $x + \frac{f_t(w_t) + \varepsilon^2}{x}$ to EWOO to obtain $\lambda_{t+1}\sqrt{m}$.

Setting $\varepsilon = 1/\sqrt[4]{T}$ gives a guarantee of form $\tilde{O}((D^3 + \min\{1/v^*, \sqrt[4]{T}\})\sqrt{T})$. Note this rate might be improvable by using the fact that v is lower-bounded due to the γ-constraint; however, we do not focus on this since this component is not the dominant term in the regret. There also exist algorithms with slightly worse $\tilde{O}(T^{\frac{3}{5}})$ regret for a simpler FTL-based algorithm for this setting [27].

3.3 Putting the two together

Now that we have an algorithm for both the initialization and the step-size, we can combine the two in Algorithm 3 to meta-learn the parameter of the exponential forecaster. Then we can obtain a bound on the task-averaged regret from Lemma 3.1 to attain our final result.

Theorem 3.2. *Define* $B^2 = D_{KL}(\mathbf{w}^* || \hat{\mathbf{v}})$, $G^2 = \frac{1}{T}\sum_{t=1}^{T} \frac{1}{(\mathrm{vol}(C_t))^2}$, *and* $D^2 \geq \frac{1}{T}\sum_{t=1}^{T} \log \frac{1}{\mathrm{vol}(C_t)} = O(\beta d \log m)$. *Then Algorithm 3 with parameters set as in Theorem 3.1 and Corollary 3.1 yields task-averaged regret*

$$\tilde{O}\left(\min\left\{\frac{\sqrt{BG}}{V\sqrt[4]{T}}, \frac{\sqrt[4]{BG}}{\sqrt[8]{T}}\right\} + 2V\right)\sqrt{m} + g(m) \tag{10}$$

Here V is the task-similarity (4).

So as in past work in meta-learning, this achieves the goal of adapting to the task-similarity by attaining asymptotic regret of $2V\sqrt{m} + O(m^{-\beta})$ on-average, where here we substitute the dispersion term for g and V^2 is the task-similarity encoding the average probability mass assigned to the different task balls by the optimal initialization distribution. We include the minimum of two rates in the bound, with the rate being $1/\sqrt[4]{T}$ is the task-similarity is a constant $\Theta_T(1)$ and $1/\sqrt[8]{T}$ if it is extremely small. As discussed in above, this rate reflects the difficulty of our meta-problem, in which we are optimizing non-smooth functions over a space of distributions; in contrast, past meta-update procedures have taken advantage of nice properties of Bregman divergences to obtain faster rates [27].

4 Applications to multi-task clustering and robustness

Meta-Learning to Cluster, an empirical study. We design experiments to evaluate our new meta-initialization algorithm for clustering on real and simulated data. Our experiments show the usefulness of our techniques in learning a sequence of piecewise-Lipschitz functions. We consider the parameterized α-Llyod's algorithm family introduced in [14]. The performance of the algorithm is measured using the Hamming loss relative to the optimal clustering, and is a piecewise constant function of α. We can compute the pieces of this function for $\alpha \in [0, 10]$ by iteratively computing the subset of parameter values where a candidate point can be the next center.

We use the small split of the *Omniglot* dataset [30], and create clustering tasks by drawing random samples consisting of five characters each, where four characters are constant throughout. We also create a Gaussian mixture binary classification dataset where each class is a 2D Gaussian distribution consisting of 100 points each, with variance $\begin{pmatrix} \sigma & 0 \\ 0 & 2\sigma \end{pmatrix}$ and centers $(0, 0)$ and $(d\sigma, 0)$. We pick $d \in [2, 3]$ to create different tasks. For both datasets we learn using 30 instances each of 10 training tasks and evaluate average loss over 5 test tasks. We perform 100 iterations to average over the

Table 1: Effect of meta-initialization on few-shot learning of clustering parameters. Performance is computed as a fraction of the average Hamming accuracy of the offline optimum parameter.

Dataset	Omniglot		Gaussian Mixture	
	One-shot	Five-shot	One-shot	Five-shot
Single task	$88.67 \pm 0.47\%$	$95.02 \pm 0.19\%$	$90.10 \pm 1.10\%$	$91.43 \pm 0.44\%$
Meta-initialized	$89.65 \pm 0.49\%$	$96.05 \pm 0.15\%$	$95.76 \pm 0.60\%$	$96.39 \pm 0.27\%$

randomization of the clustering algorithm and the exponential forecaster algorithm. We perform meta-initialization with parameters $\gamma = \eta = 0.01$ (no hyperparameter search performed). The step-size is set to minimize the regret term in Theorem 2.1, and not meta-learned.

Our results in Table 1 show that meta-learning an initialization, i.e. a distribution over α, for the exponential forecaster in this setting yields improved performance on both datasets, in both the one-shot and five-shot settings. The benefit of meta-learning is most pronounced for the Gaussian mixture case, and gains for Omniglot may increase with more tasks (c.f. the appendix).

Robust Online Meta-learning. In online learning, we seek to minimize a sequence of loss functions, and are required to perform well relative to the optimal choice in hindsight. It is possible for the observed loss functions to be noisy on some inputs, either naturally or due to adversarial intent. We will now explore the conditions under which learning robust to such an adversarial influence (i.e. outlier injection) is possible, which is particularly common in meta-learning with diverse sources.

Setup: At round i, we play x_i, observe perturbed loss $\tilde{l}_i : \mathcal{X} \to [0, 1]$ which is set by the adversary by modifying the true loss $l_i : \mathcal{X} \to [0, 1]$ using an *attack function* $a_i : \mathcal{X} \to [0, 1]$ such that $\tilde{l}_i = l_i + a_i$ and may be non-Lipschitz, and suffer perturbed loss $\tilde{l}_i(x_i)$ and true loss $l_i(x_i)$. We seek to minimize regret relative to best fixed action in hindsight, i.e. $\tilde{R}_m = \sum_{i=1}^m \tilde{l}_i(x_i) - \min_{x \in \mathcal{X}} \sum_{i=1}^m \tilde{l}_i(x)$ for the perturbed loss and regret $R_m = \sum_{i=1}^m l_i(x_i) - \min_{x \in \mathcal{X}} \sum_{i=1}^m l_i(x)$ for the true loss.

No regret can be achieved provided the adversary distribution is sufficiently smooth, i.e. satisfies β-dispersion for some $\beta > 0$, as this corresponds to online optimization of the perturbed loss function. We can show this for both perturbed and true loss. The perturbed loss guarantee is immediate from standard results on online learning of piecewise Lipschitz functions [10, 9]. For the true loss, we can achieve no regret if the adversary perturbation a_i is limited to small balls and the centers of the balls are dispersed, which we capture using the following definition.

Definition 4.1 (δ-bounded, β_a-dispersed attack)**.** *An attack function a_i is δ-bounded if there exists a ball $\mathcal{B}(x_a, \delta)$ of radius δ such that $a_i(x) = 0$ for each $x \in \mathcal{X} \setminus \mathcal{B}(x_a, \delta)$. x_a is called a center c_{a_i} for attack a_i. A sequence of attack functions $a_1, \ldots, a_m$ is said to be β_a-dispersed, if the positions of attack centers x_a are dispersed i.e. for all m and for all $\epsilon \geq m^{-\beta_a}$, $\mathbb{E}\left[\max_{x, x' \in \mathcal{X}, x \in \mathcal{B}(x', \epsilon)} \left|\{i \in [m] \mid x = c_{a_i}\}\right|\right] \leq \tilde{O}(\epsilon m)$.*

Theorem 4.1. *Given a sequence of β-dispersed adversarially perturbed losses $\tilde{l}_i = l_i + a_i$, where $\tilde{l}_i, l_i, a_i$ are piecewise L-Lipschitz functions $\mathcal{X} \to [0, 1]$ for $i = 1, \ldots, m$ and $\mathcal{X} \subset \mathbb{R}^d$, the exponential forecaster algorithm has $\mathbb{E}[\tilde{R}_m] = \tilde{O}(m\lambda + \frac{\log(1/Z)}{\lambda} + (L+1)m^{1-\beta})$ (with Z as in Theorem 2.1). If in addition we have that a_i is a $m^{-\beta_a}$-bounded, β_a-dispersed attack, then $\mathbb{E}[R_m] = \tilde{O}(m\lambda + \frac{\log(1/Z)}{\lambda} + (L+1)m^{1-\min\{\beta, \beta_a\}})$.*

Together with Theorem 3.2, this implies no regret meta-learning in the presence of dispersed adversaries, in particular the occurrence of unreliable data in small dispersed parts of the domain.

5 Conclusion

In this paper we studied the initialization-based meta-learning of piecewise-Lipschitz functions, demonstrating how online convex optimization over an adaptive discretization can find an initialization that improves the performance of the exponential forecaster across tasks, assuming the tasks have related optima. We then applied this result in two settings: online configuration of clustering algorithms and adversarial robustness in online learning. For the latter we introduced a dispersion-based understanding of robustness that we believe to be of independent interest. In addition, there are further interesting applications of our work to other algorithm configuration problems.

References

[1] Naman Agarwal, Brian Bullins, Elad Hazan, Sham Kakade, and Karan Singh. Online control with adversarial disturbances. In *International Conference on Machine Learning*, pages 111–119. PMLR, 2019.

[2] Pierre Alquier, The Tien Mai, and Massimiliano Pontil. Regret bounds for lifelong learning. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.

[3] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020.

[4] Maria-Florina Balcan. Book chapter Data-Driven Algorithm Design. In *Beyond Worst Case Analysis of Algorithms, T. Roughgarden (Ed)*. Cambridge University Press, 2020.

[5] Maria-Florina Balcan, Avrim Blum, Dravyansh Sharma, and Hongyang Zhang. On the power of abstention and data-driven decision making for adversarial robustness. *arXiv preprint arXiv:2010.06154*, 2020.

[6] Maria-Florina Balcan, Avrim Blum, and Santosh Vempala. Efficient representations for lifelong learning and autoencoding. In *Proceedings of the 28th Annual Conference on Learning Theory*, 2015.

[7] Maria-Florina Balcan, Travis Dick, and Manuel Lang. Learning to link. In *International Conference on Learning Representations*, 2019.

[8] Maria-Florina Balcan, Travis Dick, Tuomas Sandholm, and Ellen Vitercik. Learning to branch. In *International conference on machine learning*, pages 344–353. PMLR, 2018.

[9] Maria-Florina Balcan, Travis Dick, and Dravyansh Sharma. Learning piecewise Lipschitz functions in changing environments. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pages 3567–3577, 2020.

[10] Maria-Florina Balcan, Travis Dick, and Ellen Vitercik. Dispersion for data-driven algorithm design, online learning, and private optimization. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 603–614, 2018.

[11] Maria-Florina Balcan, Vaishnavh Nagarajan, Ellen Vitercik, and Colin White. Learning-theoretic foundations of algorithm configuration for combinatorial partitioning problems. In *Annual Conference on Learning Theory*, pages 213–274, 2017.

[12] Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. A general theory of sample complexity for multi-item profit maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 173–174, 2018.

[13] Maria-Florina Balcan and Dravyansh Sharma. Data driven algorithms for limited labeled data learning. *arXiv preprint arXiv:2103.10547*, 2021.

[14] Maria-Florina F Balcan, Travis Dick, and Colin White. Data-driven clustering via parameterized lloyd's families. *Advances in Neural Information Processing Systems*, 31:10641–10651, 2018.

[15] Avrim Blum. Technical perspective: Algorithm selection as a learning problem. *Communications of the ACM*, 63(6):86–86, 2020.

[16] Wieland Brendel, Jonas Rauber, Alexey Kurakin, Nicolas Papernot, Behar Veliqi, Sharada P Mohanty, Florian Laurent, Marcel Salathé, Matthias Bethge, Yaodong Yu, et al. Adversarial vision challenge. In *The NeurIPS'18 Competition*, pages 129–153. Springer, 2020.

[17] Fei Chen, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. Federated meta-learning for recommendation. arXiv, 2018.

[18] Vincent Cohen-Addad and Varun Kanade. Online optimization of smoothed piecewise constant functions. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 2017.

[19] Giulia Denevi, Carlo Ciliberto, Riccardo Grazzi, and Massimiliano Pontil. Learning-to-learn stochastic gradient descent with biased regularization. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

[20] Giulia Denevi, Carlo Ciliberto, Riccardo Grazzi, and Massimiliano Pontil. Online-within-online meta-learning. In *Advances in Neural Information Processing Systems*, 2019.

[21] Simon S. Du, Wei Hu, Sham M. Kakade, Jason D. Lee, and Qi Lei. Few-shot learning via learning the representation, provably.

[22] Yan Duan, Marcin Andrychowicz, Bradly Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in Neural Information Processing Systems*, 2017.

[23] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. On the convergence theory of gradient-based model-agnostic meta-learning algorithms. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2020.

[24] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.

[25] Rishi Gupta and Tim Roughgarden. A PAC approach to application-specific algorithm selection. *SIAM Journal on Computing*, 46(3):992–1017, 2017.

[26] Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69:169–192, 2007.

[27] Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. Adaptive gradient-based meta-learning methods. In *Advances in Neural Information Processing Systems*, 2019.

[28] Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. Provable guarantees for gradient-based meta-learning. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.

[29] Weihao Kong, Raghav Somani, Sham Kakade, and Sewoong Oh. Robust meta-learning for mixed linear regression with small batches. *Advances in Neural Information Processing Systems*, 33, 2020.

[30] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015.

[31] Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(1):2853–2884, 2016.

[32] Anh Nguyen, Jason Yosinski, and Jeff Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 427–436, 2015.

[33] Francesco Orabona, Nicolo Cesa-Bianchi, and Claudio Gentile. Beyond logarithmic bounds in online learning. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, 2012.

[34] Francesco Orabona and David Pal. Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems*, 2016.

[35] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *arXiv preprint arXiv:1912.13445*, 2019.

[36] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? Towards understanding the effectiveness of MAML. In *Proceedings of the 8th International Conference on Learning Representations*, 2020.

[37] Alon Resler and Yishay Mansour. Adversarial online learning with noise. In *International Conference on Machine Learning*, pages 5429–5437. PMLR, 2019.

[38] Nikunj Saunshi, Yi Zhang, Mikhail Khodak, and Sanjeev Arora. A sample complexity separation between non-convex and convex meta-learning. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.

[39] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.

[40] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.

[41] Sebastian Thrun and Lorien Pratt. *Learning to Learn*. Springer Science & Business Media, 1998.

[42] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *European Symposium on Research in Computer Security*, pages 480–501. Springer, 2020.

[43] Nilesh Tripuraneni, Chi Jin, and Michael I. Jordan. Provable meta-learning of linear representations. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.

[44] Yufan Zhou, Zhenyi Wang, Jiayi Xian, Changyou Chen, and Jinhui Xu. Meta-learning with neural tangent kernels. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Alongside contributions in context, e.g. the end of Section 3.1.
   (c) Did you discuss any potential negative societal impacts of your work? [No] Our concern w.r.t. the negative societal impact of this theoretical work is limited to standard risks associated with ML, e.g. for privacy or fair treatment.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A] Experiments run on personal computer (16GB, 2.3 GHz Dual-Core).
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]

(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]