

Prediction in Online Convex Optimization for Parametrizable Objective Functions

Robert Ravier¹ and Vahid Tarokh¹

¹Department of Electrical and Computer Engineering, Duke University

Abstract—Many techniques for online optimization problems involve making decisions based solely on presently available information: fewer works take advantage of potential predictions. In this paper, we discuss the problem of online convex optimization for parametrizable objectives, i.e. optimization problems that depend solely on the value of a parameter at a given time. We introduce a new regularity for dynamic regret based on the accuracy of predicted values of the parameters and show that, under mild assumptions, accurate prediction can yield tighter bounds on dynamic regret. Inspired by recent advances on learning how to optimize, we also propose a novel algorithm to simultaneously predict and optimize for parametrizable objectives and study its performance using simulated and real data.

I. INTRODUCTION

Online convex optimization (OCO) has received significant attention in recent years due to its wide range of applicability. These applications include ad selection [13], video streaming [16], power scheduling [22] among many others. We refer the reader to [13], [24] for a more rigorous introduction, but we briefly summarize below.

The typical OCO scenario can be modeled as the following game. At time $t - 1$, the player must pick a candidate point x_t for belonging to some constraint set \mathcal{X} . At time t , the true convex loss function $f_t(\cdot)$ is revealed, and the player suffers a loss $f_t(x_t)$. This continues for a total of T time steps. The general goal is to find an algorithm that performs well with respect to some notion of regret. Historically, much focus has been given towards the static regret, i.e. performance with respect to the optimal fixed point in hindsight:

$$\mathbf{Reg}_S(\{x_t\}) := \sum_{t=1}^T f_t(x_t) - \min_{x^* \in \mathcal{X}} \left(\sum_{t=1}^T f_t(x^*) \right) \quad (1)$$

where $\{x_t\}$ is the sequence of moves played. It has been well-established that algorithms can achieve sublinear [29] and even logarithmic regret under suitable assumptions [12]. The same cannot be said of the measure of regret of interest in our paper, dynamic regret (sometimes called *restricted dynamic regret* [27] or *tracking regret* [11]). Dynamic regret measures the performance with respect to the optimal values of the function at each time, i.e.:

$$\mathbf{Reg}_D(\{x_t\}) := \sum_{t=1}^T \left(f_t(x_t) - \min_{x_t^* \in \mathcal{X}} f_t(x_t^*) \right) \quad (2)$$

OCO algorithms have been more frequently analyzed by looking at their dynamic regret [15], [18], [21], [26], [28].

The analysis of these algorithms does not focus on sublinear regret (as this is impossible to achieve in general [26]) but rather focuses on bounding performance in terms of different regularities depending on the specific algorithm employed.

Somewhat surprisingly, few of the algorithms for OCO explicitly make use of predictions of future objective functions and gradients. Indeed, OCO has historically been viewed from an adversarial lens. This is perhaps too conservative for many scenarios. For example, in applications related to power allocation, frequently past information concerning usage is indicative of the future. Leveraging accurate predictions could better assist in many scenarios where online optimization techniques are utilized. However, the effect of accurate predictions of relevant data points on the performance of such online algorithms is generally not clear. Some algorithms give bounds on regret in terms of the accuracy of blackbox predictions of functions or gradients [15], [20], but it is not immediately obvious as to how one can get these from data. Other methods assume that accurate predictions are known throughout the duration of the scenario (e.g. [18]), or assume very particular structure of the resulting accuracy or potentially involve unrealistic assumptions such as fully optimizing a function at every point and time (e.g. [8]). We would like to address these issues in our paper for a large class of objective functions.

To this end, we make the observation that in many situations, the specific *form* of the objective function is known and fixed throughout time. More specifically:

Definition 1. Let $f(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ be a function that is convex in the first argument. An optimization problem is **parametric** if it is of the form

$$\min_{x \in \mathcal{X}} f(x, \theta)$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ is a closed convex set for some θ in a parameter space $\Theta \subset \mathbb{R}^m$.

This is a general form of optimization problem present in predictive optimization problems [8], [14]. Throughout the remainder of the paper, we assume that we are interested in OCO problems where the cost functions f_t are of the form $f_t(x) = f(x, \theta_t)$. We note that many important objective functions of theoretical and practical interest are encompassed by this specific form.

Example 1. Let $g_1(x), \dots, g_m(x)$ be convex functions with

domain $X \subseteq \mathbb{R}^n$ and let $\Theta = \Delta^m$ be the standard unit m -simplex. Then

$$f(x, \theta_t) := \theta_t^1 g_1(x) + \cdots + \theta_t^m g_m(x)$$

is a functional time series.

Example 2. For a collection of assets A , let μ_t and Σ_t be their corresponding sample mean and covariance, and let $\lambda_t > 0$. For $X = \Delta^{|A|}$ the unit simplex on the number of assets $|A|$, the Markowitz optimal portfolio [19] with respect to $[\mu_t, \Sigma_t, \lambda_t]$ is the argmin over X of the following function

$$f(x, [\mu_t, \Sigma_t, \lambda_t]) = x^T \Sigma_t x - \lambda_t x^T \mu_t \quad (3)$$

For parametric optimization problems, prediction of objective functions and relevant quantities reduces to prediction of parameters, a much more well-studied though still difficult problem (e.g. [4], [6]).

The main results of our paper are summarized as follows.

- We show that, under mild regularity assumptions, gradient descent using predicted objectives for parametric optimization problems as defined in Definition 1 can improve the dynamic regret over standard online gradient descent provided sufficient accuracy in predicted values. The method of proof of our dynamic regret bounds is general enough to extend to cases of where a descent algorithm yields a contraction, i.e. we have some inequality of the form

$$\|x_{t+1} - x_{t+1}^*\| \leq \rho \|x_t - x_t^*\| \quad (4)$$

for some $0 < \rho < 1$.

- We provide a meta-learning algorithm called SMAD, inspired by recent innovations in learning how to optimize, that simultaneously learns the optimal parameter prediction process from a collection of models while performing descent.

The remainder of the paper is summarized as follows. In Section II we further detail preliminary details and assumptions needed for the remainder of the paper. In Section III we detail our theoretical results concerning the performance of predictive online gradient descent. In Section IV we detail our meta-learning algorithm for simultaneous modeling and descent. In Section V we give numerical simulations to backup our intuition and evaluate our algorithm's performance. In Section VI we make concluding remarks.

II. PRELIMINARIES

A. Regularities for Dynamic Regret

As discussed earlier, dynamic regret bounds for algorithms focus on various regularities of the OCO problem of interest. These regularities generally focus not on algorithmic decisions but on properties of the elements of the problem outside of the algorithm's control. We briefly a number of these quantities. One of the more prevalent regularities, notably appearing in [21], [29] is the path length of the sequence of optimal points: if $x_t^* = \operatorname{argmin}_{x \in X} f_t(x)$ for f_t convex, then

$$\mathcal{P}^* := \sum_{t=1}^{T-1} \|x_t^* - x_{t+1}^*\|$$

Other regularities of interest include the *squared* path length introduced in [28], functional variation [2] and the gradient variation [9], which are measurements that depend on the sup norm of the differences between the functions and their gradients between times $t-1$ and t .

More directly relevant to our discussion are what we term *prediction regularities*. These are not regularities as above in the sense that they are within an algorithm's purview. Nevertheless, they have demonstrated importance for certain regret bounds. [15] introduces a squared predictive gradient regularity

$$\sum_{t=1}^T \|M_t - \nabla_x f_t(x_t)\|^2$$

where M_t is a prediction of $\nabla_x f_t(x_t)$ prior to it becoming revealed. We will consider predictive regularities consistent with the parametric optimization framework we outlined earlier. For $f_t(x) = f(x, \theta_t)$, we consider the *parameter prediction regularity*

$$P^\theta := \sum_{t=2}^T \|\theta_t - \hat{\theta}_t\|$$

This quantity measures cumulative prediction error respect to the parameters (hence objective functions) and will be important to our subsequent analysis. One can also similarly introduced a squared parameter prediction regularity in a manner analogous to the squared path length, obtained by squaring the norms in the above term, but we do not pursue this.

B. Theoretical Assumptions

We now detail the theoretical assumptions needed for the remainder of the paper. The first few assumptions are standard for dynamic regret analysis in OCO and can be seen in, for example, [21], [28]. Recall that all functions we consider will be of the form $f_t(x) = f(x, \theta_t)$ and that our closed, convex constraint set is given by X with corresponding projection Π_X and the parameter set is given by Θ .

Assumption 1. The function $f(x, \theta)$ is Lipschitz continuous in x , i.e. there exists a constant $G > 0$ such that, for all $x, y \in X$ and $\theta \in \Theta$

$$|f(x, \theta) - f(y, \theta)| \leq G \|x - y\|.$$

Assumption 2. The function $f(x, \theta)$ is L -smooth in x , i.e. there exists a constant $G > 0$ such that

The function $f(x, \theta)$ is λ -strongly convex in x , i.e. there exists a constant $G > 0$ such that

$$f(y, \theta) \leq f(x, \theta) + \nabla_x f(x, \theta)^T (y - x) + \frac{L}{2} \|y - x\|^2$$

Algorithm 1 Online Predictive Gradient Descent

Input: Step size $\eta > 0$, and $x_1 \in \mathcal{X}$
for $t = 1 \rightarrow T$ **do**
 Receive parameter θ_t
 Predict $\hat{\theta}_{t+1}$ from $\theta_t, \dots, \theta_1$
 Compute $x_{t+1} = \Pi_{\mathcal{X}}(x_t - \eta \nabla_x f(x, \hat{\theta}_{t+1}))$
end for

Assumption 3. The function $f(x, \theta)$ is λ -strongly convex in x , i.e. there exists a constant $G > 0$ such that

$$f(y, \theta) \geq f(x, \theta) + \nabla_x f(x, \theta)^T (y - x) + \frac{\lambda}{2} \|y - x\|^2$$

The first two assumptions are common throughout the OCO literature and give upper bounds on the first and second derivatives. Assumption 3 is a more recent assumption in the OCO literature, first appearing in dynamic regret analysis in [21], but is a common assumption for the analysis of descent algorithms like gradient descent [5]. Quadratic functions defined over a compact set are examples of functions satisfying all three assumptions.

For our analysis, we will need an additional regularity assumption concerning the behavior of gradients with respect to θ .

Assumption 4. The function $f(x, \theta)$ has Lipschitz continuous x -gradients in θ , i.e. there exists some $C_\theta > 0$ such that, for all $\theta_1, \theta_2 \in \Theta$ and $x \in \mathcal{X}$, we have

$$\|\nabla_x f(x, \theta_1) - \nabla_x f(x, \theta_2)\| \leq C_\theta \|\theta_1 - \theta_2\|.$$

It is not hard to check that functional time series (Example 1) satisfies Assumption 4 provided that the sum of the $\nabla_x(g_1(x) + \dots + g_m(x))$ is bounded. It is similarly easy to see that the Markowitz portfolio function in Example 2 also satisfies Assumption 4 when viewing the collection of parameters as a column-stacked vector.

III. THEORETICAL RESULTS FOR PREDICTION IN DESCENT

We now analyze gradient descent when incorporating prediction. See Algorithm 1 for the pseudocode. The main difference between Algorithm 1 and standard online gradient descent is in the prediction step in the form of the parameter prediction. As prediction can mean many different things depending on the situation at hand, we avoid mentioning a particular process at this stage.

In analyzing Algorithm 1, we are specifically interested in bounding the dynamic regret in terms of the path-length expressions \mathcal{P}^* and \mathcal{S}^* as well as the parameter prediction regularity P^θ . To this end, we have the following result.

Theorem 1. Let Assumptions 1-4 hold. If $\eta \leq 1/L$, then for a $C_{\eta, \lambda} < 1$ we have the following bound on regret of Algorithm 1:

$$\mathbf{Reg}_D(\{x_t\}) \leq \frac{G\|x_1 - x_1^*\|}{1 - C_{\eta, \lambda}} + \frac{GC_{\eta, \lambda}}{1 - C_{\eta, \lambda}} \mathcal{P}^* + \frac{G\eta C_\theta}{1 - C_{\eta, \lambda}} P^\theta \quad (5)$$

The following lemma, which we state without proof from [28], makes the constant $C_{\eta, \lambda}$ in the above theorem more precise.

Lemma 1. Let $g(x)$ be a λ -strongly convex function and L smooth with minimum attained at x^* . Then projected gradient descent is a contraction provided that $\eta \leq \frac{1}{L}$: if \mathcal{X} is the constraint set and $\Pi_{\mathcal{X}}$ is the projection onto \mathcal{X} , we have

$$\|\Pi_{\mathcal{X}}(v - \eta \nabla g(v)) - x^*\| \leq C_{\lambda, \eta} \|v - x^*\|$$

$$\text{where } C_{\lambda, \eta} := \sqrt{1 - \frac{2\lambda\eta}{1 + \eta\lambda}}.$$

The proof of the theorem is similar to other calculations of dynamic regret with a few modifications to accommodate the difference in descent strategy. We briefly summarize the proof methodology and defer the exact details to the Appendix, though the interested reader will also find the discussion in the proof of Corollary 1 enlightening. Previous bounds on the regret for online gradient descent were built around the fact that the descent direction at time t for guessing x_{t+1} was $\nabla_x f_t$ and not $\nabla_x f_{t+1}$. However, we are explicitly trying to descend using a prediction of f_{t+1} . Our error in this end will be driven by the quality of our prediction of $\nabla_x f_{t+1}$. Assumption 4 gives that this is controlled by the quality of the parameter prediction. Standard bounding and rearranging then gives the proof of the theorem.

We discuss the results of the bound. When comparing our result to the results of [21] and [28], we notice that an additional multiplicative constant less than 1 appears in front of the path length \mathcal{P}^* at the cost of the entire P^θ term. If we have perfect prediction, or even near perfect, this allows us to achieve a smaller regret bound, potentially significantly so depending on the exact quantity of \mathcal{P}^* . For imperfect prediction, there is a tradeoff, and it is possible that previous regret bounds are superior in some instances. We detail some examples.

Example 3. Assume that the actual amount of time between θ_t and θ_{t+1} is Δt . Assume that we know that θ_t satisfies some ordinary differential equation $\dot{\theta}_t = V(\theta_t, t)$. Runge-Kutta methods can be used to numerically integrate the ODE and compute a predicted value $\hat{\theta}_{t+1}$ whose error is of the order $(\Delta t)^4$. Sufficiently small values of Δt will thus guarantee small contribution from the prediction regularity.

Example 4. Following [8], we consider the case that our predictions are accurate up to noise, i.e. $\theta_t - \hat{\theta}_t = h_1^t X_1^t + \dots + h_k^t X_k^t$ for each t where the X_i^t are i.i.d. mean zero sub-Gaussian random variables with variance parameter σ^2 and the h_i^t are constants. Though this implies that P^θ is a random variable, it is well known that such a linear combination satisfies a high probability bound:

$$\mathbb{P}\left(\left|\sum_{i=1}^k h_i^t X_i^t\right| > \varepsilon\right) \leq \exp\left(-\frac{\varepsilon^2}{\sigma^2 \left(\sum_{i=1}^k (h_i^t)^2\right)}\right)$$

This implies that P^θ will also be small with high probability.

We now consider the case where we may wish to perform multiple descent steps between each iteration. This will inject

powers of $C_{\eta,\lambda}$ into the above regret bound, potentially increasing the error bound. Surprisingly, this does not significantly affect the constant in front of P^θ .

Corollary 1. *Let Assumptions 1-4 hold. Consider Algorithm 1 with the single gradient descent step replaced by k gradient descent steps. If $\eta \leq 1/L$, then we have the following bound on regret of this revised version Algorithm 1:*

$$\text{Reg}_D(\{x_t\}) \leq \frac{G\|x_1 - x_1^*\|}{1 - C_{\eta,\lambda}^k} + \frac{GC_{\eta,\lambda}^k}{1 - C_{\eta,\lambda}^k} \mathcal{P}^* + \frac{G\eta C_\theta}{1 - C_{\eta,\lambda}} P^\theta \quad (6)$$

Proof. The same idea as the single gradient proof holds, but the estimate is slightly different. Indeed, we can replace the one step gradient descent with a k -step gradient to also get a contraction, with $C_{\eta,\lambda}$ being replaced by $C_{\eta,\lambda}^k$. However, as we are computing a gradient descent step with $\hat{\theta}_t$ and not θ_t , we must use a triangle inequality at every iteration before we can use the contractive estimate.

More precisely, let z_t^k be the point obtained by using k predicted gradient descent steps from x_{t-1} . We can repeat the above analysis to see

$$\begin{aligned} \|z_t^k - x_t^*\| &= \|z_t^{k-1} - \eta \nabla_x f(z_t^{k-1}, \hat{\theta}_t) - x_t^*\| \\ &\leq \|z_t^{k-1} - \eta \nabla_x f(z_t^{k-1}, \theta_t) - x_t^*\| \\ &\quad + \eta \|\nabla_x f(z_t^{k-1}, \theta_t) - \nabla_x f(z_t^{k-1}, \hat{\theta}_t)\| \\ &\leq C_{\eta,\lambda} \|z_t^{k-1} - x_t^*\| + \eta C_\theta \|\theta_t - \hat{\theta}_t\| \end{aligned}$$

A routine induction gives

$$\begin{aligned} \|z_t^k - x_t^*\| &\leq C_{\eta,\lambda}^k \|x_{t-1} - x_t^*\| + C_\theta \eta \|\theta_t - \hat{\theta}_t\| \sum_{i=0}^{k-1} C_{\eta,\lambda}^i \\ &\leq C_{\eta,\lambda}^k \|x_{t-1} - x_t^*\| + \frac{C_\theta \eta (1 - C_{\eta,\lambda}^k)}{1 - C_{\eta,\lambda}} \|\theta_t - \hat{\theta}_t\| \end{aligned}$$

where the last inequality follows by majorizing the summation in the first inequality by an infinite series and subsequent summation. \square

We end this section by investigating the generalizability of our proof to other descent methods. The main idea outlined above is to replace the estimated descent step with the true descent step and then estimate the error by looking at both the contraction of the true descent direction and the prediction error of the descent step. More formally, we consider a general descent algorithm, where the gradient descent step is replaced by $x_{t+1} = \Pi_{\mathcal{X}}(x_t - \eta D_{\text{est}})$ for some descent direction D_{est} . If D_{true} is the descent direction for the true value of the parameter and D_{est} is that for the predicted value of the parameter, then for this general version, we have the following estimate:

$$\begin{aligned} \|\Pi_{\mathcal{X}}(x_t - \eta D_{\text{est}}) - x_{t+1}^*\| &\leq \|\Pi_{\mathcal{X}}(x_t - \eta D_{\text{true}}) - x_{t+1}^*\| \\ &\quad + \|\Pi_{\mathcal{X}}(x_t - \eta D_{\text{true}}) - \Pi_{\mathcal{X}}(x_t - \eta D_{\text{est}})\| \\ &\leq \|\Pi_{\mathcal{X}}(x_t - \eta D_{\text{true}}) - x_{t+1}^*\| \\ &\quad + \|\eta D_{\text{true}} - \eta D_{\text{est}}\| \end{aligned}$$

by nonexpansiveness of the projection. Provided that the true descent direction yields a contraction, then with sufficient regularity of the descent directions with respect to the parameter θ , we can get an analogous result for different regularity assumptions on $f(x, \theta)$. For example, if we were to pick a Newton step, so $D_{\text{true}} = (\nabla_x^2 f(x, \theta_{t+1}))^{-1} \nabla_x f(x, \theta_{t+1})$ and similarly for D_{est} , we can bound the difference by

$$\begin{aligned} \|D_{\text{true}} - D_{\text{est}}\| &\leq \|A^{-1} (\nabla_x f(x, \theta_{t+1}) - \nabla_x f(x, \hat{\theta}_{t+1}))\| \\ &\quad + \left\| \left((\nabla_x^2 f(x, \theta_{t+1}))^{-1} - (\nabla_x^2 f(x, \hat{\theta}_{t+1}))^{-1} \right) B \right\| \end{aligned}$$

where $A = \nabla_x^2 f(x, \theta_{t+1})$ and $B = \nabla_x f(x, \hat{\theta}_{t+1})$. Under sufficient regularity of the operator norm of the inverse Hessian of f , we can obtain a similar expression in terms of P^θ as we did for gradient descent in Theorem 1. We do not investigate this further, but instead use the above to illustrate that our technique is not particular to the gradient descent step used.

IV. SMAD: SIMULTANEOUS MODELING AND DESCENT

The theoretical results in the previous section are rather general, and give regret bounds in terms of the quality of prediction without specifying how this prediction is done in general. Frequently we do not know the true process generating the data we are trying to predict, but instead have a collection of candidate models for which we hope at least one will make accurate predictions. To this end, we would like to develop a practical algorithm that will gradually learn the best data generating model over which to optimize among a predetermined collection of models.

In particular, we follow the ideas of MetaGrad and Ader [25], [27] and employ an approach based on expert learning. Expert learning has been well studied (see, e.g., [7]) and is summarized as follows. Each expert corresponds to a particular class of data-generating models (for example, each expert can correspond to the lag of an autoregressive (AR) model). We assume that each expert knows the specific function $f(x, \theta)$ that we are trying to optimize. At time $t-1$, each expert makes a prediction $\hat{\theta}_t^i$ as to the future value of θ and uses this prediction in order to evaluate its own predictive online gradient descent procedure, thus giving a predicted value of x denoted x_t^i . Each expert suffers a loss based on evaluating $f(x_t^i, \theta_t)$, the experts are reweighted using a Gibbs posterior update procedure, and the process is repeated until the last time point is reached.

We briefly comment on the conditional statement in our expert descent algorithm. From a practical standpoint, when one wishes to start modeling data with a collection of models, they may not have a sufficient amount of data to reliably use a particular model. For example, if one wishes to model a time series of data with an AR(k) process via the Yule-Walker equations, one cannot reliably estimate a model if the number of data points is not sufficiently large relative to k . To accommodate for this, we allow the user to specify whether/when additional models are added, and do this by mixing a new model in by reweighting the predictive distribution. When a new model comes online, we propose to initialize its candidate x -value at the previous point output by the algorithm. We will

Algorithm 2 Expert Learning Algorithm for Simultaneous Descent and Data Modeling

Input: Parameters $\beta, \eta, \gamma > 0$, initial data $\theta_{-i_0}, \dots, \theta_0$
 Output: $p_t = [p_{t,1}, \dots, p_{t,N}]$ (predictive distribution over the collection of N models), x_t (predictive points for optimization)
 Initialize $w_{1,0} = \dots = w_{k_0,0} = 1$, $w_{k,0} = 0$ for $k > k_0$
 Initialize $M = k_0$
for $t = 1 \rightarrow T$ **do**
 if Initializing new model **then**
 Compute $w_{i,t-1} = (1 - \beta)w_{i,t-1}$ for $i \leq M$
 Compute $w_{M+1,t-1} = \beta$
 Compute $p_{i,t-1} = \frac{w_{i,t-1}}{\sum_{i=1}^N w_{i,t-1}}$
 Initialize $x_{t-1}^{M+1} = x_{t-1}$
 Compute $M = M + 1$
 end if
 Observe θ_t
 Initialize $x_t = 0$
 for $1 \leq i \leq M$ **do**
 Receive $v_i = \Pi_X(x_{t-1}^i - \eta \nabla_x f(x_{t-1}^i, \hat{\theta}_t^i))$ from expert i
 Compute $x_t = x_t + p_{i,t-1} v_i$
 Compute $l_i = \exp[-\gamma f(X(v_i), \theta_t)]$
 Compute $w_{i,t} = p_{i,t-1} l_i$
 end for
 Output $x_t = \Pi_X(x_t)$
 Compute $w_{i,t} = p_{i,t-1} l_i$
 Compute $p_{i,t} = \frac{w_{i,t}}{\sum_{i=1}^N w_{i,t}}$
end for

make use of these procedures in our numerical examples, but for ease of presentation avoid this in our theoretical analysis.

We present the results of our theoretical analysis below.

Theorem 2. *Assume Assumptions 1-4 hold and that the range of $f(x, \theta)$, $D := \sup_{x,\theta} f(x, \theta) - \inf_{x,\theta} f(x, \theta)$ is bounded. Then if Algorithm (2) introduces no additional models upon starting, it has regret bound*

$$\begin{aligned} \text{Reg}_D(\{x_t\}) &\leq \frac{G\|x_1 - x_1^*\|}{1 - C_{\eta,\lambda}} + \frac{GC_{\eta,\lambda}}{1 - C_{\eta,\lambda}} \mathcal{P}^* \\ &+ \frac{G\eta C_\theta}{1 - C_{\eta,\lambda}} \min_i P_i^\theta + \frac{D\sqrt{2T}}{4} (1 + \ln N) \end{aligned}$$

where P_i^θ is the parameter prediction regularity for model i and \min_i is with respect to the available models.

The proof combines standard techniques for evaluating expert learning algorithms as in [7] along with previous regret analysis for predictive online gradient descent and is reserved for the Appendix.

V. NUMERICAL EXPERIMENTS

We now detail numerical experiments investigating the efficacy of prediction in OCO for parametric objectives as well as the performance of our objective function. We will primarily compare our prediction related results to standard online gradient descent (OGD), where the descent direction is

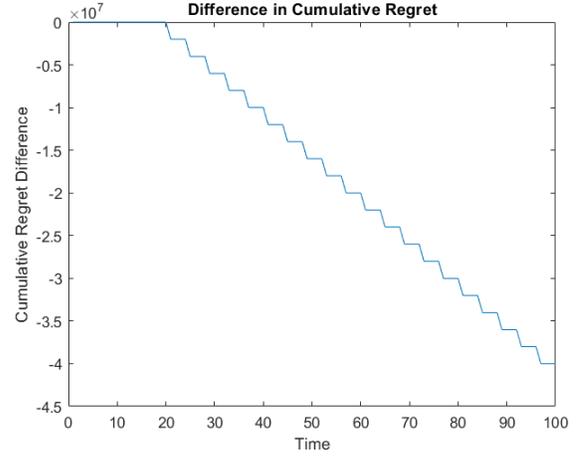


Fig. 1: Difference in dynamic regret between Predictive OGD and OGD. Lower numbers indicate better performance for Predictive OGD.

fully determined by the value of the objective function at the present. From an intuitive perspective, if the process governing the objective function is relatively stationary with small variation between time steps, we would anticipate minimal difference between standard OGD and the predictive version we laid out above. The main differences should arise when there are predictable but significant jumps in the parameter governing the objective function for which a method with close to accurate models will predict reasonably well, whereas OGD will suffer a loss for not catching the jump.

A. OGD versus Prediction with Fixed Model

We consider the case where we have one reasonable candidate model of the objective function parameter. This first experiment is adapted directly from [21]. We consider the parametric objective function

$$\begin{aligned} f_t([x_1, x_2]) &= f([x_1, x_2], [a_t, b_t, c_t]) \\ &= 100(x_1 - a_t)^2 + (x_2 - b_t)^2 + c_t \end{aligned}$$

where $[a_t, b_t, c_t]$ alternates between $[-100, 0, 30] + \epsilon_t$ and $[100, 20, -50] + \epsilon_t$ every four iterations, where ϵ_t is three-dimensional Gaussian noise with mean zero and covariance $10I_3$. The constraint set is the disc centered at the origin with radius 50. We compare the performance of OGD with the performance of the following procedure: follow OGD for the first 10 timesteps, then estimate a two dimensional AR(4) model for $[a_t, b_t]$ using the Yule-Walker equations, and use predictive online gradient descent. Both methods will be initialized at $[0, 40]$. Following the convention in [21], we set the step size for both methods to be $1/200$.

The results of the experiment averaged over fifty repetitions can be found in Figure 1. The main validation measure we employ is the difference in cumulative regret between the predictive version of OGD and the standard version of OGD. For this measure, lower values indicate better performance of the predictive method. As expected, the curve remains flat

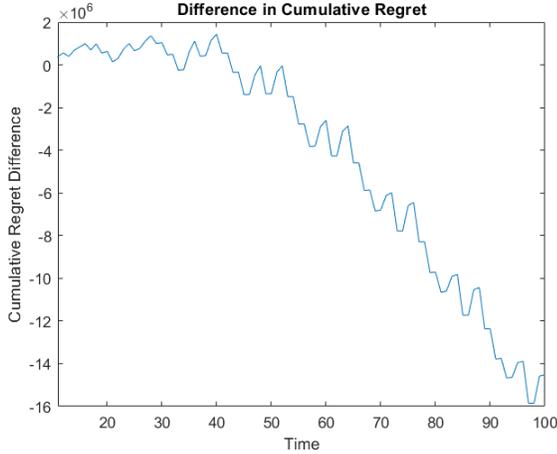


Fig. 2: Difference in dynamic regret between our expert learning method and OGD. Lower numbers indicate better performance for the expert learning method.

for the first 10 timesteps as the descent method is the same. When the model estimation turns starts, the predictive method begins to outperform the standard OGD as evidenced by the gradually decreasing curve in the figure. The steps on the curve is indicative of the step-like behavior of the cumulative OGD regret as observed in [21].

B. Expert Learning on Synthetic Data

We would now like to test out Algorithm 2 in a misspecified setting, i.e. when the model classes do not contain the true objective parameter generating process. To this end, we keep most of the same settings as in the first experiment, but we change the switching process: $[a_t, b_t, c_t]$ alternates between $[-100, 0, 30] + \epsilon_t$ for four time steps and $[100, 20, -50] + \epsilon_t$ for six time steps. The models that we use are AR(k) for k between 1 to 5. We first observe the process for ten time steps before using an AR(1), and then add a new AR model every ten time steps until all five are active. All models are again estimated by the Yule Walker equations. For the other parameters of the algorithm, we use $\beta = 0.2$ and $\gamma = 5 \times 10^{-7}$.

The results of the experiment averaged over fifty repetitions can be found in Figure 2. Initially, the algorithm performs worse than OGD, which is not surprising as the initial models are not close to the parameter process. As higher lag AR models are added and the available models include ones closer to the true process, the performance of the expert learning method improves, eventually becoming the clear favorite over standard OGD.

C. Expert Learning on Financial Data

As a real world application, we now consider the problem of portfolio optimization. Though there are a number of different ways to construct portfolios in both online and offline settings (see, e.g., [10], [19], [23]), we restrict our attention to Markowitz portfolio theory as introduced in Example 2. To remind the reader, given a collection of assets z_1, \dots, z_n , the Markowitz optimal portfolio allocation is given by:

$$\operatorname{argmin}_{x \in \mathbb{R}^n} x^T \Sigma_t x - \lambda_t x^T \mu_t \quad (7)$$

$$x_1, \dots, x_n \geq 0, \sum_{i=1}^n x_i = 1$$

where Σ represents the covariance matrix for the returns of the assets, μ is the average return of each asset, and $\lambda > 0$ is a parameter encoding the tradeoff between expected returns and risk; optimizing with $\lambda = 0$ finds the portfolio with the least amount of risk.

In this framework, we consider a hypothetical scenario of a client with a rapidly changing, noisy risk tolerance. In this scenario, a portfolio manager reaches out to his client every month (30 days) with a series of new portfolios, each of which is constructed built by estimation of their client’s risk tolerance and subsequent optimizing of Equation (3) given different lookback periods on a given collection of assets to be between 15 and 90 days in increments of 15 for computing relevant means and covariances. The risk tolerance is estimated by a series of autoregressive processes on the risk tolerance with lags between 30 and 180 days in increments of 30¹.

Unbeknownst to the manager, the client evaluates the portfolio also via Equation (3), but with a 50 day lookback period and a risk generated by the following process. For the first 240 days, the client’s risk parameter is $\max(4 + \epsilon, 0)$, where ϵ is Gaussian noise of mean zero and variance 0.64. The remaining risk parameters are generated as follows. Setting $b_0 = 4$, we have $\lambda_t = \max(b_t + \epsilon_t, 0)$, where b_t satisfies for $t \geq 1$:

$$b_{t+1} = \begin{cases} b_t & \text{with probability } 0.9 \\ \text{Unif}(1, \dots, 20) & \text{with probability } 0.1 \end{cases}$$

and the ϵ_t are Gaussian with mean zero and variance 0.64. Here, $\text{Unif}(1, \dots, 20)$ denotes the discrete uniform distribution on the integers 1, ..., 20. As data, we make use of the NYSE dataset used frequently in the portfolio optimization literature, a collection of 36 stock returns taken over a period of 22 years [3], [17]. We also add a risk-free asset that gives constant, low returns of 1% every 360 days compounded daily. The goal, as with the synthetic experiment, is to predict the best portfolio of those offered that optimizes the client’s objective function without knowing that objective function in the future.

Initially, both methods started with portfolio uniform across the assets. Both method used a descent step size of 0.1 and a learning rate of $\gamma = 50$. The projection step is performed by setting negative amounts of assets equal to zero before normalizing the percentages of remaining assets. We assume that the learner has 10 months to observe the risk trends before starting, which gives a sufficient amount of data in order to estimate the risk process via the Yule Walker equations. This renders the mixing constant β irrelevant to this simulation. Our evaluation period occurs over 150 months.

The results of this experiment averaged over 200 times can be seen in Figure 3. Unsurprisingly, there is plenty of noise in the resulting curve due to the potentially large fluxuations in

¹Risk is only observed every 30 days

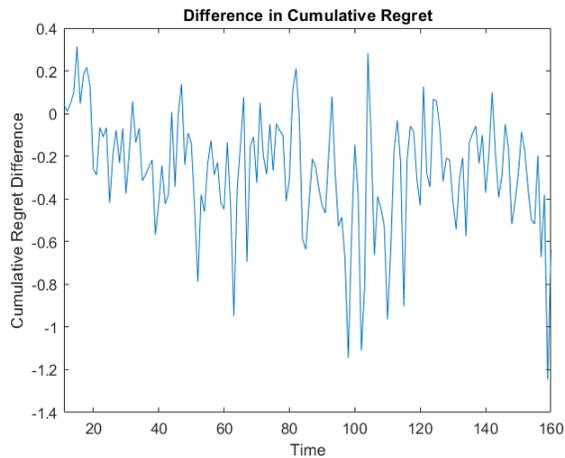


Fig. 3: Difference in dynamic regret between our expert learning method and OGD. Lower numbers indicate better performance for the expert learning method.

the risk parameter. Nevertheless, though there are occasionally spikes for which standard OGD does better, we notice that the difference in cumulative regret between the expert learning and the OGD methods tends to be negative, indicating better performance by our expert learning algorithm.

VI. CONCLUSION

We have discussed the problem of online convex optimization for a wide class of parametric objective functions, for which prediction of parameters subsequently gives us predictions of objective functions. We analyzed a predictive version of online gradient descent and showed that its dynamic regret can improve on currently known bounds provided that prediction of parameters is accurate. We also proposed SMAD, an expert learning-based algorithm that allows us to simultaneously model the parameter process and optimize. We finally showed via numerical examples the power of prediction in OCO, especially in environments where sharp changes can occur, and showed that SMAD can offer better performance than standard online gradient descent in both synthetic and real data settings.

There are a number of directions in which to extend this work. It would be interesting to consider the effect of other smoothness conditions on objective functions, such as self-concordance and semi-strong convexity to see how much improvement we can get on regret bounds as was done in [28]. It would also be interesting to investigate the effect of prediction in optimization when trying to predict parametric constraint functions, though this will inevitably be challenging due to potential constraint violations from inaccurate predictions. Finally, since predictions can sometimes yield confidence intervals, it would be extremely interesting to explore this problem from the lens of robust optimization, where one focuses on minimizing the maximal possible loss. [1]

VII. ACKNOWLEDGEMENTS

This work was partially funded by DARPA grant number FA8650-18-1-7837 .

REFERENCES

- [1] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust optimization*, volume 28. Princeton University Press, 2009.
- [2] O. Besbes, Y. Gur, and A. Zeevi. Non-stationary stochastic optimization. *Operations research*, 63(5):1227–1244, 2015.
- [3] A. Borodin, R. El-Yaniv, and V. Gogan. Can we learn to beat the best stock. In *Advances in Neural Information Processing Systems*, pages 345–352, 2004.
- [4] G. E. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [5] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [6] P. J. Brockwell, R. A. Davis, and M. V. Calder. *Introduction to time series and forecasting*, volume 2. Springer, 2002.
- [7] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- [8] N. Chen, J. Comden, Z. Liu, A. Gandhi, and A. Wierman. Using predictions in online optimization: Looking forward with an eye on the past. *ACM SIGMETRICS Performance Evaluation Review*, 44(1):193–206, 2016.
- [9] C.-K. Chiang, T. Yang, C.-J. Lee, M. Mahdavi, C.-J. Lu, R. Jin, and S. Zhu. Online optimization with gradual variations. In *Conference on Learning Theory*, pages 6–1, 2012.
- [10] T. M. Cover. Universal portfolios. In *The Kelly Capital Growth Investment Criterion: Theory and Practice*, pages 181–209. World Scientific, 2011.
- [11] E. Hall and R. Willett. Dynamical models and tracking regret in online convex programming. In *International Conference on Machine Learning*, pages 579–587, 2013.
- [12] E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- [13] E. Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- [14] S. Ito, A. Yabe, and R. Fujimaki. Unbiased objective estimation in predictive optimization. In *International Conference on Machine Learning*, pages 2181–2190, 2018.
- [15] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan. Online optimization: Competing with dynamic comparators. In *Artificial Intelligence and Statistics*, pages 398–406, 2015.
- [16] V. Joseph and G. de Veciana. Jointly optimizing multi-user rate adaptation for video transport over wireless systems: Mean-fairness-variability tradeoffs. In *INFOCOM, 2012 Proceedings IEEE*, pages 567–575. IEEE, 2012.
- [17] B. Li and S. C. Hoi. Online portfolio selection: A survey. *ACM Computing Surveys (CSUR)*, 46(3):35, 2014.
- [18] Y. Li, G. Qu, and N. Li. Using predictions in online optimization with switching costs: A fast algorithm and a fundamental limit. In *2018 Annual American Control Conference (ACC)*, pages 3008–3013. IEEE, 2018.
- [19] H. Markowitz. Portfolio selection. *The journal of finance*, 7(1):77–91, 1952.
- [20] M. Mohri and S. Yang. Accelerating online convex optimization via adaptive prediction. In *Artificial Intelligence and Statistics*, pages 848–856, 2016.
- [21] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro. Online optimization in dynamic environments: Improved regret rates for strongly convex problems. In *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pages 7195–7201. IEEE, 2016.
- [22] B. Narayanaswamy, V. K. Garg, and T. Jayram. Online optimization for the smart (micro) grid. In *Proceedings of the 3rd international conference on future energy systems: where energy, computing and communication meet*, page 19. ACM, 2012.
- [23] A. Ponsich, A. L. Jaimes, and C. A. C. Coello. A survey on multiobjective evolutionary algorithms for the solution of the portfolio optimization problem and other finance and economics applications. *IEEE Transactions on Evolutionary Computation*, 17(3):321–344, 2013.
- [24] S. Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.
- [25] T. van Erven and W. M. Koolen. Metagrad: Multiple learning rates in online learning. In *Advances in Neural Information Processing Systems*, pages 3666–3674, 2016.
- [26] T. Yang, L. Zhang, R. Jin, and J. Yi. Tracking slowly moving clairvoyant: Optimal dynamic regret of online learning with true and noisy gradient. In *International Conference on Machine Learning*, pages 449–457, 2016.

- [27] L. Zhang, S. Lu, and Z.-H. Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems*, pages 1330–1340, 2018.
- [28] L. Zhang, T. Yang, J. Yi, J. Rong, and Z.-H. Zhou. Improved dynamic regret for non-degenerate functions. In *Advances in Neural Information Processing Systems*, pages 732–741, 2017.
- [29] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 928–936, 2003.

APPENDIX

A. Proof of Theorem 1

The dynamic regret upper bound computation is similar to others. Recall that x_t^* denotes the minimizers. From the assumption that x -gradients are bounded above by a constant G for all θ , we have

$$\sum_{t=1}^T f(x_t, \theta_t) - f(x_t^*, \theta_t) \leq G \sum_{t=1}^T \|x_t - x_t^*\| = G\|x_1 - x_1^*\| + G \sum_{t=2}^T \|x_t - x_t^*\| \quad (8)$$

To bound the sum on the right-hand side, observe by the definition of the algorithm and the triangle inequality, we have:

$$\begin{aligned} \sum_{t=2}^T \|x_t - x_t^*\| &= \sum_{t=2}^T \|\Pi_X(x_{t-1} - \eta \nabla_x f(x_{t-1}, \hat{\theta}_t) - x_t^*)\| \\ &\leq \sum_{t=2}^T \left(\|\Pi_X(x_{t-1} - \eta \nabla_x f(x_{t-1}, \theta_t)) - x_t^*\| + \eta \|\nabla_x f(x_{t-1}, \theta_t) - \nabla_x f(x_{t-1}, \hat{\theta}_t)\| \right) \end{aligned}$$

We implicitly used the fact that the projection operator is nonexpansive in our setting in the second inequality. So, we have, referencing Lemma 1:

$$\begin{aligned} \sum_{t=2}^T \|x_t - x_t^*\| &\leq \sum_{t=2}^T \left(\|\Pi_X(x_{t-1} - \eta \nabla_x f(x_{t-1}, \theta_t) - x_t^*)\| + \eta \|\nabla_x f(x_{t-1}, \theta_t) - \nabla_x f(x_{t-1}, \hat{\theta}_t)\| \right) \\ &\leq \sum_{t=2}^T \left(C_{\eta, \lambda} \|x_{t-1} - x_t^*\| + \eta \|\nabla_x f(x_{t-1}, \theta_t) - \nabla_x f(x_{t-1}, \hat{\theta}_t)\| \right) \\ &\leq \sum_{t=2}^T \left(C_{\eta, \lambda} \|x_{t-1} - x_t^*\| + \eta C_\theta \|\theta_t - \hat{\theta}_t\| \right) \\ &\leq C_{\eta, \lambda} \sum_{t=2}^T \|x_{t-1} - x_{t-1}^*\| + C_{\eta, \lambda} \sum_{t=2}^T \|x_{t-1}^* - x_t^*\| + \eta C_\theta \sum_{t=2}^T \|\theta_t - \hat{\theta}_t\| \\ &\leq C_{\eta, \lambda} \|x_1 - x_1^*\| + C_{\eta, \lambda} \sum_{t=2}^T \|x_t - x_t^*\| + C_{\eta, \lambda} \sum_{t=2}^T \|x_{t-1}^* - x_t^*\| + \eta C_\theta \sum_{t=2}^T \|\theta_t - \hat{\theta}_t\| \end{aligned}$$

where the second inequality follows from Lemma 1, the third inequality follows from the Lipschitz continuity in θ of the x -gradients, the fourth inequality from the triangle inequality, and the last inequality from the norm being nonnegative. Subtracting the second term on the last line from both sides, because $C_{\eta, \lambda} < 1$, we see that

$$\sum_{t=2}^T \|x_t - x_t^*\| \leq \frac{C_{\eta, \lambda}}{1 - C_{\eta, \lambda}} \left(\|x_1 - x_1^*\| + \sum_{t=1}^{T-1} \|x_t^* - x_{t+1}^*\| \right) + \frac{C_\theta \eta}{1 - C_{\eta, \lambda}} \sum_{t=2}^T \|\theta_t - \hat{\theta}_t\| \quad (9)$$

Plugging this estimate into Equation 8 and some minor algebra completes the proof.

B. Proof of Theorem 2

The basic idea of the proof is to follow the basic regret calculation for the expert learning algorithm, and then use the result of Theorem 1. The expert learning calculation is routine and adapted from [7] for the sake of completeness. Assume that the models are indexed from 1 to N . Then for model i , define $L_t^i = \sum_{k=1}^t f_k(x_k^i)$ and

$$W_t = \sum_{i=1}^N w_{i,0} e^{-\gamma L_t^i}$$

From this, it follows that properties of the logarithm that

$$\ln W_T \geq \left(\ln \max_{1 \leq i \leq N} w_{i,0} e^{-\gamma L_T^i} \right) = -\gamma \left(\min_{1 \leq i \leq N} \left(L_T^i + \frac{1}{\gamma} \ln \frac{1}{w_{i,0}} \right) \right) = -\gamma \left(\min_{1 \leq i \leq N} \left(L_T^i + \frac{1}{\gamma} \ln N \right) \right) \quad (10)$$

By logarithm properties, we also have that $\ln W_T = \ln W_1 + \sum_{t=1}^{T-1} \ln \left(\frac{W_t}{W_{t-1}} \right)$. It is not hard to see by the definition of $w_{i,t}$, we have

$$\ln\left(\frac{W_t}{W_{t-1}}\right) = \ln\left(\sum_{i=1}^N w_{i,t} e^{-\gamma f_t(x_t^i)}\right) \quad (11)$$

Note that the sum on the left is an expectation of the random variable $e^{-\gamma f_t(x_t^i)}$, so Hoeffding's and Jensen's inequalities gives

$$\ln\left(\frac{W_t}{W_{t-1}}\right) \leq -\gamma\left(\sum_{i=1}^N w_{i,t} f_t(x_t^i)\right) + \frac{\gamma^2 D^2}{8} \leq -\gamma f_t\left(\sum_{i=1}^N w_{i,t} x_t^i\right) + \frac{\gamma^2 D^2}{8} \leq -\gamma f(x_t) + \frac{\gamma^2 D^2}{8} \quad (12)$$

Combining Equations 10, 11, and 12, we have

$$-\gamma\left(\min_{1 \leq i \leq N} \left(L_T^i + \frac{1}{\gamma} \ln N\right)\right) \leq \ln W_T \leq -\gamma \sum_{t=1}^T f_t(x_t) + \frac{\gamma^2 T D^2}{8} \quad (13)$$

Cancelling a γ on both sides and rearranging Equation 13 gives

$$\sum_{t=1}^T f_t(x_t) - \min_i \left(\sum_{t=1}^T f_t(x_t^i)\right) \leq \frac{T \gamma D^2}{8} - \frac{1}{\gamma} \ln N \quad (14)$$

Routine calculus minimizes the right hand side of Equation 14 by setting $\gamma = \sqrt{8/(TD^2)}$ to get

$$\sum_{t=1}^T f_t(x_t) - \left(\sum_{t=1}^T f_t(x_t^i)\right) \leq \frac{D\sqrt{2T}}{4}(1 + \ln N) \quad (15)$$

for every i . Since each model follows its own version of predictive online gradient descent, each model has its own bound according to Theorem 1, namely:

$$\sum_{t=1}^T f_t(x_t^i) - \left(\sum_{t=1}^T f_t(x_t^*)\right) \leq \frac{G\|x_1 - x_1^*\|}{1 - C_{\eta,\lambda}} + \frac{GC_{\eta,\lambda}}{1 - C_{\eta,\lambda}} \mathcal{P}^* + \frac{G\eta C_\theta}{1 - C_{\eta,\lambda}} \min_i P_i^\theta \quad (16)$$

Combining Equations 15 and 16 for the i with the smallest bound for Equation 16 completes the proof.