
Online Convex Optimization with Unconstrained Domains and Losses

Ashok Cutkosky
Department of Computer Science
Stanford University
ashokc@cs.stanford.edu

Kwabena Boahen
Department of Bioengineering
Stanford University
boahen@stanford.edu

Abstract

We propose an online convex optimization algorithm (RESCALED EXP) that achieves optimal regret in the unconstrained setting without prior knowledge of any bounds on the loss functions. We prove a lower bound showing an *exponential* separation between the regret of existing algorithms that require a known bound on the loss functions and any algorithm that does not require such knowledge. RESCALED EXP matches this lower bound asymptotically in the number of iterations. RESCALED EXP is naturally hyperparameter-free and we demonstrate empirically that it matches prior optimization algorithms that require hyperparameter optimization.

1 Online Convex Optimization

Online Convex Optimization (OCO) [1, 2] provides an elegant framework for modeling noisy, antagonistic or changing environments. The problem can be stated formally with the help of the following definitions:

Convex Set: A set W is convex if W is contained in some real vector space and $tw + (1-t)w' \in W$ for all $w, w' \in W$ and $t \in [0, 1]$.

Convex Function: $f : W \rightarrow \mathbb{R}$ is a convex function if $f(tw + (1-t)w') \leq tf(w) + (1-t)f(w')$ for all $w, w' \in W$ and $t \in [0, 1]$.

An OCO problem is a game of repeated rounds in which on round t a learner first chooses an element w_t in some convex space W , then receives a convex loss function ℓ_t , and suffers loss $\ell_t(w_t)$. The *regret* of the learner with respect to some other $u \in W$ is defined by

$$R_T(u) = \sum_{t=1}^T \ell_t(w_t) - \ell_t(u)$$

The objective is to design an algorithm that can achieve low regret with respect to any u , even in the face of adversarially chosen ℓ_t .

Many practical problems can be formulated as OCO problems. For example, the stochastic optimization problems found widely throughout machine learning have exactly the same form, but with i.i.d. loss functions, a subset of the OCO problems. In this setting the goal is to identify a vector w_* with low generalization error ($\mathbb{E}[\ell(w_*) - \ell(u)]$). We can solve this by running an OCO algorithm for T rounds and setting w_* to be the average value of w_t . By online-to-batch conversion results [3, 4], the generalization error is bounded by the expectation of the regret over the ℓ_t divided by T . Thus, OCO algorithms can be used to solve stochastic optimization problems while also performing well in non-i.i.d. settings.

The regret of an OCO problem is upper-bounded by the regret on a corresponding Online Linear Optimization (OLO) problem, in which each ℓ_t is further constrained to be a *linear* function: $\ell_t(w) = g_t \cdot w_t$ for some g_t . The reduction follows, with the help of one more definition:

Subgradient: $g \in W$ is a subgradient of f at w , denoted $g \in \partial f(w)$, if and only if $f(w) + g \cdot (w' - w) \leq f(w')$ for all w' . Note that $\partial f(w) \neq \emptyset$ if f is convex.¹

To reduce OCO to OLO, suppose $g_t \in \partial \ell_t(w_t)$, and consider replacing $\ell_t(w)$ with the linear approximation $g_t \cdot w$. Then using the definition of subgradient,

$$R_T(u) = \sum_{t=1}^T \ell_t(w_t) - \ell_t(u) \leq \sum_{t=1}^T g_t(w_t - u) = \sum_{t=1}^T g_t w_t - g_t u$$

so that replacing $\ell_t(w)$ with $g_t \cdot w$ can only make the problem more difficult. All of the analysis in this paper therefore addresses OLO, accessing convex losses functions only through subgradients.

There are two major factors that influence the regret of OLO algorithms: the size of the space W and the size of the subgradients g_t . When W is a bounded set (the “constrained” case), then given $B = \max_{w \in W} \|w\|$, there exist OLO algorithms [5, 6] that can achieve $R_T(u) \leq O\left(BL_{\max}\sqrt{T}\right)$ without knowing $L_{\max} = \max_t \|g_t\|$. When W is unbounded (the “unconstrained” case), then given L_{\max} , there exist algorithms [7, 8, 9] that achieve $R_T(u) \leq \tilde{O}(\|u\| \log(\|u\|)L_{\max}\sqrt{T})$ or $R_t(u) \leq \tilde{O}(\|u\| \sqrt{\log(\|u\|)}L_{\max}\sqrt{T})$, where \tilde{O} hides factors that depend logarithmically on L_{\max} and T . These algorithms are known to be optimal (up to constants) for their respective regimes [10, 7]. All algorithms for the unconstrained setting to-date require knowledge of L_{\max} to achieve these optimal bounds.² Thus a natural question is: *can we achieve $O(\|u\| \log(\|u\|))$ regret in the unconstrained, unknown- L_{\max} setting?* This problem has been posed as a COLT 2016 open problem [12], and is solved in this paper.

A simple approach is to maintain an estimate of L_{\max} and double it whenever we see a new g_t that violates the assumed bound (the so-called “doubling trick”), thereby turning a known- L_{\max} algorithm into an unknown- L_{\max} algorithm. This strategy fails for previous known- L_{\max} algorithms because their analysis makes strong use of the assumption that each and every $\|g_t\|$ is bounded by L_{\max} . The existence of even a small number of bound-violating g_t can throw off the entire analysis.

In this paper, we prove that it is actually impossible to achieve regret $O\left(\|u\| \log(\|u\|)L_{\max}\sqrt{T} + L_{\max} \exp\left[\left(\max_t \frac{\|g_t\|}{L(t)}\right)^{1/2-\epsilon}\right]\right)$ for any $\epsilon > 0$ where L_{\max} and $L(t) = \max_{t' < t} \|g_{t'}\|$ are unknown in advance (Section 2). This immediately rules out the “ideal” bound of $\tilde{O}(\|u\| \sqrt{\log(\|u\|)}L_{\max}\sqrt{T})$ which is possible in the known- L_{\max} case. Secondly, we provide an algorithm, RESCALEDEXP, that matches our lower bound without prior knowledge of L_{\max} , leading to a naturally hyperparameter-free algorithm (Section 3). To our knowledge, this is the first algorithm to address the unknown- L_{\max} issue while maintaining $O(\|u\| \log \|u\|)$ dependence on u . Finally, we present empirical results showing that RESCALEDEXP performs well in practice (Section 4).

2 Lower Bound with Unknown L_{\max}

The following theorem rules out algorithms that achieve regret $O(u \log(u)L_{\max}\sqrt{T})$ without prior knowledge of L_{\max} . In fact, any such algorithm must pay an up-front penalty that is *exponential* in T . This lower bound resolves a COLT 2016 open problem (Parameter-Free and Scale-Free Online Algorithms) [12] in the negative.

¹In full generality, a subgradient is an element of the dual space W^* . However, we will only consider cases where the subgradient is naturally identified with an element in the original space W (e.g. W is finite dimensional) so that the definition in terms of dot-products suffices.

²There are algorithms that do not require L_{\max} , but achieve only regret $O(\|u\|^2)$ [11]

Theorem 1. For any constants $c, k, \epsilon > 0$, there exists a T and an adversarial strategy picking $g_t \in \mathbb{R}$ in response to $w_t \in \mathbb{R}$ such that regret is:

$$\begin{aligned} R_T(u) &= \sum_{t=1}^T g_t w_t - g_t u \\ &\geq (k + c\|u\| \log \|u\|) L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \exp((2T)^{1/2-\epsilon}) \\ &\geq (k + c\|u\| \log \|u\|) L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \exp \left[\left(\max_t \frac{\|g_t\|}{L(t)} \right)^{1/2-\epsilon} \right] \end{aligned}$$

for some $u \in \mathbb{R}$ where $L_{\max} = \max_{t \leq T} \|g_t\|$ and $L(t) = \max_{t' < t} \|g_{t'}\|$.

Proof. We prove the theorem by showing that for sufficiently large T , the adversary can “checkmate” the learner by presenting it only with the subgradient $g_t = -1$. If the learner fails to have w_t increase quickly, then there is a $u \gg 1$ against which the learner has high regret. On the other hand, if the learner ever does make w_t higher than a particular threshold, the adversary immediately punishes the learner with a subgradient $g_t = 2T$, again resulting in high regret.

Let T be large enough such that both of the following hold:

$$\frac{T}{4} \exp\left(\frac{T^{1/2}}{4 \log(2)c}\right) > k \log(2) \sqrt{T} + k \exp((2T)^{1/2-\epsilon}) \quad (1)$$

$$\frac{T}{2} \exp\left(\frac{T^{1/2}}{4 \log(2)c}\right) > 2kT \exp((2T)^{1/2-\epsilon}) + 2kT \sqrt{T} \log(2T + 1) \quad (2)$$

The adversary plays the following strategy: for all $t \leq T$, so long as $w_t < \frac{1}{2} \exp(T^{1/2}/4 \log(2)c)$, give $g_t = -1$. As soon as $w_t \geq \frac{1}{2} \exp(T^{1/2}/4 \log(2)c)$, give $g_t = 2T$ and $g_t = 0$ for all subsequent t . Let’s analyze the regret at time T in these two cases.

Case 1: $w_t < \frac{1}{2} \exp(T^{1/2}/4 \log(2)c)$ for all t :

In this case, let $u = \exp(T^{1/2}/4 \log(2)c)$. Then $L_{\max} = 1$, $\max_t \frac{\|g_t\|}{L(t)} = 1$, and using (1) the learner’s regret is at least

$$\begin{aligned} R_T(u) &\geq Tu - T \frac{1}{2} \exp\left(\frac{T^{1/2}}{4 \log(2)c}\right) \\ &= \frac{1}{2} Tu \\ &= cu \log(u) \sqrt{T} \log(2) + \frac{T}{4} \exp\left(\frac{T^{1/2}}{4 \log(2)c}\right) \\ &> cu \log(u) L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \exp((2T)^{1/2-\epsilon}) \\ &= (k + cu \log u) L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \exp \left[(2T)^{1/2-\epsilon} \right] \end{aligned}$$

Case 2: $w_t \geq \frac{1}{2} \exp(T^{1/2}/4 \log(2)c)$ for some t :

In this case, $L_{\max} = 2T$ and $\max_t \frac{\|g_t\|}{L(t)} = 2T$. For $u = 0$, using (2), the regret is at least

$$\begin{aligned} R_T(u) &\geq \frac{T}{2} \exp\left(\frac{T^{1/2}}{4 \log(2)c}\right) \\ &\geq 2kT \exp((2T)^{1/2-\epsilon}) + 2kT \sqrt{T} \log(2T + 1) \\ &= k L_{\max} \exp((2T)^{1/2-\epsilon}) + k L_{\max} \sqrt{T} \log(L_{\max} + 1) \\ &= (k + cu \log u) L_{\max} \sqrt{T} \log(L_{\max} + 1) + k L_{\max} \exp \left[(2T)^{1/2-\epsilon} \right] \end{aligned}$$

□

The exponential lower-bound arises because the learner has to move exponentially fast in order to deal with exponentially far away u , but then experiences exponential regret if the adversary provides a gradient of unprecedented magnitude in the opposite direction. However, if we play against an adversary that is constrained to give loss vectors $\|g_t\| \leq L_{\max}$ for some L_{\max} that does not grow with time, or if the losses do not grow too quickly, then we can still achieve $R_T(u) = O(\|u\| \log(\|u\|) L_{\max} \sqrt{T})$ asymptotically without knowing L_{\max} . In the following sections we describe an algorithm that accomplishes this.

3 RESCALEDEXP

Our algorithm, RESCALEDEXP, adapts to the unknown L_{\max} using a guess-and-double strategy that is robust to a small number of bound-violating g_t s. We initialize a guess L for L_{\max} to $\|g_1\|$. Then we run a novel known- L_{\max} algorithm that can achieve good regret in the unconstrained u setting. As soon as we see a g_t with $\|g_t\| > 2L$, we update our guess to $\|g_t\|$ and restart the known- L_{\max} algorithm. To prove that this scheme is effective, we show (Lemma 3) that our known- L_{\max} algorithm does not suffer too much regret when it sees a g_t that violates its assumed bound.

Our known- L_{\max} algorithm uses the Follow-the-Regularized-Leader (FTRL) framework. FTRL is an intuitive way to design OCO algorithms [13]: Given functions $\psi_t : W \rightarrow \mathbb{R}$, at time T we play $w_T = \operatorname{argmin} \left[\psi_{T-1}(w) + \sum_{t=1}^{T-1} \ell_t(w) \right]$. The functions ψ_t are called regularizers. A large number of OCO algorithms (e.g. gradient descent) can be cleanly formulated as instances of this framework.

Our known- L_{\max} algorithm is FTRL with regularizers $\psi_t(w) = \psi(w)/\eta_t$, where $\psi(w) = (\|w\| + 1) \log(\|w\| + 1) - \|w\|$ and η_t is a scale-factor that we adapt over time. Specifically, we set $\eta_t^{-1} = k\sqrt{2}\sqrt{M_t + \|g\|_{1:t}^2}$, where we use the compressed sum notations $g_{1:T} = \sum_{t=1}^T g_t$ and $\|g\|_{1:T}^2 = \sum_{t=1}^T \|g_t\|^2$. M_t is defined recursively by $M_0 = 0$ and $M_t = \max(M_{t-1}, \|g_{1:t}\|/p - \|g\|_{1:t}^2)$, so that $M_t \geq M_{t-1}$, and $M_t + \|g\|_{1:t}^2 \geq \|g_{1:t}\|/p$. k and p are constants: $k = \sqrt{2}$ and $p = L_{\max}^{-1}$.

RESCALEDEXP's strategy is to maintain an estimate L_t of L_{\max} at all time steps. Whenever it observes $\|g_t\| \geq 2L_t$, it updates $L_{t+1} = \|g_t\|$. We call periods during which L_t is constant *epochs*. Every time it updates L_t , it restarts our known- L_{\max} algorithm with $p = \frac{1}{L_t}$, beginning a new epoch. Notice that since L_t at least doubles every epoch, there will be at most $\log_2(L_{\max}/L_1) + 1$ total epochs. To address edge cases, we set $w_t = 0$ until we suffer a non-constant loss function, and we set the initial value of L_t to be the first non-zero g_t . Pseudo-code is given in Algorithm 1, and Theorem 2 states our regret bound. For simplicity, we re-index so that that g_1 is the first non-zero gradient received. No regret is suffered when $g_t = 0$ so this does not affect our analysis.

Algorithm 1 RESCALEDEXP

Initialize: $k \leftarrow \sqrt{2}$, $M_0 \leftarrow 0$, $w_1 \leftarrow 0$, $t_* \leftarrow 1$ // t_* is the start-time of the current epoch.
for $t = 1$ **to** T **do**
 Play w_t , receive subgradient $g_t \in \partial \ell_t(w_t)$.
 if $t = 1$ **then**
 $L_1 \leftarrow \|g_1\|$
 $p \leftarrow 1/L_1$
 end if
 $M_t \leftarrow \max(M_{t-1}, \|g_{t_*:t}\|/p - \|g\|_{t_*:t}^2)$.
 $\eta_t \leftarrow \frac{1}{k\sqrt{2(M_t + \|g\|_{t_*:t}^2)}}$
 //Set w_{t+1} using FTRL update
 $w_{t+1} \leftarrow -\frac{g_{t_*:t}}{\|g_{t_*:t}\|} [\exp(\eta_t \|g_{t_*:t}\|) - 1] = \operatorname{argmin}_w \left[\frac{\psi(w)}{\eta_t} + g_{t_*:t} w \right]$
 if $\|g_t\| > 2L_t$ **then**
 //Begin a new epoch: update L and restart FTRL
 $L_{t+1} \leftarrow \|g_t\|$
 $p \leftarrow 1/L_{t+1}$
 $t_* \leftarrow t + 1$
 $M_t \leftarrow 0$
 $w_{t+1} \leftarrow 0$
 else
 $L_{t+1} \leftarrow L_t$
 end if
end for

Theorem 2. *Let W be a separable real inner-product space with corresponding norm $\|\cdot\|$ and suppose (with mild abuse of notation) every loss function $\ell_t : W \rightarrow \mathbb{R}$ has some subgradient $g_t \in W^*$ at w_t such that $g_t(w) = g_t \cdot w$ for some $g_t \in W$. Let $M_{\max} = \max_t M_t$. Then if $L_{\max} = \max_t \|g_t\|$*

and $L(t) = \max_{t' < t} \|g_{t'}\|$, *rescaledexp* achieves regret:

$$\begin{aligned} R_T(u) &\leq (2\psi(u) + 96) \left(\log_2 \left(\frac{L_{\max}}{L_1} \right) + 1 \right) \sqrt{M_{\max} + \|g\|_{1:T}^2} \\ &\quad + 8L_{\max} \left(\log_2 \left(\frac{L_{\max}}{L_1} \right) + 1 \right) \min \left[\exp \left(8 \max_t \frac{\|g_t\|^2}{L(t)^2} \right), \exp(\sqrt{T/2}) \right] \\ &= O \left(L_{\max} \log \left(\frac{L_{\max}}{L_1} \right) \left[(\|u\| \log(\|u\|) + 2)\sqrt{T} + \exp \left(8 \max_t \frac{\|g_t\|^2}{L(t)^2} \right) \right] \right) \end{aligned}$$

The conditions on W in Theorem 2 are fairly mild. In particular they are satisfied whenever W is finite-dimensional and in most kernel method settings [14]. In the kernel method setting, W is an RKHS of functions $\mathcal{X} \rightarrow \mathbb{R}$ and our losses take the form $\ell_t(w) = \ell_t(\langle w, k_{x_t} \rangle)$ where k_{x_t} is the representing element in W of some $x_t \in \mathcal{X}$, so that $g_t = \bar{g}_t k_{x_t}$ where $\bar{g}_t \in \partial \ell_t(\langle w, k_{x_t} \rangle)$.

Although we nearly match our lower-bound exponential term of $\exp((2T)^{1/2-\epsilon})$, in order to have a practical algorithm we need to do much better. Fortunately, the $\max_t \frac{\|g_t\|^2}{L(t)^2}$ term may be significantly smaller when the losses are not fully adversarial. For example, if the loss vectors g_t satisfy $\|g_t\| = t^2$, then the exponential term in our bound reduces to a manageable constant even though $\|g_t\|$ is growing quickly without bound.

To prove Theorem 2, we bound the regret of RESCALED EXP during each epoch. Recall that during an epoch, RESCALED EXP is running FTRL with $\psi_t(w) = \psi(w)/\eta_t$. Therefore our first order of business is to analyze the regret of FTRL across one of these epochs, which we do in Lemma 3 (proved in appendix):

Lemma 3. *Set $k = \sqrt{2}$. Suppose $\|g_t\| \leq L$ for $t < T$, $1/L \leq p \leq 2/L$, $g_T \leq L_{\max}$ and $L_{\max} \geq L$. Let $W_{\max} = \max_{t \in [1, T]} \|w_t\|$. Then the regret of FTRL with regularizers $\psi_t(w) = \psi(w)/\eta_t$ is:*

$$\begin{aligned} R_T(u) &\leq \psi(u)/\eta_T + 96\sqrt{M_T + \|g\|_{1:T}^2} + 2L_{\max} \min \left[W_{\max}, 4 \exp \left(4 \frac{L_{\max}^2}{L^2} \right), \exp(\sqrt{T/2}) \right] \\ &\leq (2\psi(u) + 96) \sqrt{\sum_{t=1}^{T-1} L|g_t| + L_{\max}^2} + 8L_{\max} \min \left[\exp \left(\frac{4L_{\max}^2}{L^2} \right), \exp(\sqrt{T/2}) \right] \\ &\leq L_{\max}(2(\|u\| + 1) \log(\|u\| + 1) - \|u\|) + 96\sqrt{T} + 8L_{\max} \min \left[e^{\frac{4L_{\max}^2}{L^2}}, e^{\sqrt{T/2}} \right] \end{aligned}$$

Lemma 3 requires us to know the value of L in order to set p . However, the crucial point is that it encompasses the case in which L is misspecified on the last loss vector. This allows us to show that RESCALED EXP does not suffer too much by updating p on-the-fly.

Proof of Theorem 2. The theorem follows by applying Lemma 3 to each epoch in which L_t is constant.

Let $1 = t_1, t_2, t_3, \dots, t_n$ be the various increasing values of t_* (as defined in Algorithm 1), and we define $t_{n+1} = T + 1$. Then define

$$R_{a:b}(u) = \sum_{t=a}^{b-1} g_t(w_t - u)$$

so that $R_T(u) \leq \sum_{j=1}^n R_{t_j:t_{j+1}}(u)$. We will bound $R_{t_j:t_{j+1}}(u)$ for each j .

Fix a particular $j < n$. Then $R_{t_j:t_{j+1}}(u)$ is simply the regret of FTRL with $k = \sqrt{2}$, $p = \frac{1}{L_{t_j}}$, $\eta_t = \frac{1}{k\sqrt{2(M_t + \|g\|_{1:t}^2)}}$ and regularizers $\psi(w)/\eta_t$. By definition of L_t , for $t \in [1, t_{j+1} - 2]$ we have $\|g_t\| \leq 2L_{t_j}$. Further, if $L = \max_{t \in [1, t_{j+1} - 2]} \|g_t\|$ we have $L \geq L_{t_j}$. Therefore, $L_{t_j} \leq L \leq 2L_{t_j}$ so that $\frac{1}{L} \leq p \leq \frac{2}{L}$. Further, we have $\|g_{t_{j+1}-1}\|/L_{t_j} \leq 2 \max_t \|g_t\|/L(t)$. Thus by Lemma 3 we

have

$$\begin{aligned}
R_{t_j:t_{j+1}}(u) &\leq \psi(u)/\eta_{t_{j+1}-1} + 96\sqrt{M_{t_{j+1}-1} + \|g\|_{t_j:t_{j+1}-1}^2} \\
&\quad + 2L_{\max} \min \left[W_{\max}, 4 \exp \left(4 \frac{\|g_{t_{j+1}-1}\|^2}{L_{t_j}^2} \right), \exp \left(\frac{\sqrt{t_{j+1} - t_j}}{\sqrt{2}} \right) \right] \\
&\leq \psi(u)/\eta_{t_{j+1}-1} + 96\sqrt{M_{\max} + \|g\|_{t_j:t_{j+1}-1}^2} + 8L_{\max} \min \left[e^{8 \max_t \frac{\|g_t\|^2}{L(t)^2}}, e^{\sqrt{T/2}} \right] \\
&\leq (2\psi(u) + 96)\sqrt{M_{\max} + \|g\|_{1:T}^2} + 8L_{\max} \min \left[\exp \left(8 \max_t \frac{\|g_t\|^2}{L(t)^2} \right), \exp(\sqrt{T/2}) \right]
\end{aligned}$$

Summing across epochs, we have

$$\begin{aligned}
R_T(u) &= \sum_{j=1}^n R_{t_j:t_{j+1}}(u) \\
&\leq n \left[(2\psi(u) + 96)\sqrt{M_{\max} + \|g\|_{1:T}^2} + 8L_{\max} \min \left[\exp \left(8 \max_t \frac{\|g_t\|^2}{L(t)^2} \right), \exp(\sqrt{T/2}) \right] \right]
\end{aligned}$$

Observe that $n \leq \log_2(L_{\max}/L_1) + 1$ to prove the first line of the theorem. The big-Oh expression follows from the inequality: $M_{t_{j+1}-1} \leq L_{t_j} \sum_{t=t_j}^{t_{j+1}-1} \|g_t\| \leq L_{\max} \sum_{t=1}^T \|g_t\|$. \square

Our specific choices for k and p are somewhat arbitrary. We suspect (although we do not prove) that the preceding theorems are true for larger values of k and any p inversely proportional to L_t , albeit with differing constants. In Section 4 we perform experiments using the values for k , p and L_t described in Algorithm 1. In keeping with the spirit of designing a hyperparameter-free algorithm, no attempt was made to empirically optimize these values at any time.

4 Experiments

4.1 Linear Classification

To validate our theoretical results in practice, we evaluated RESCALEDEXP on 8 classification datasets. The data for each task was pulled from the libsvm website [15], and can be found individually in a variety of sources [16, 17, 18, 19, 20, 21, 22]. We use linear classifiers with hinge-loss for each task and we compare RESCALEDEXP to five other optimization algorithms: ADAGRAD [5], SCALEINVARIANT [23], PISTOL [24], ADAM [25], and ADADELTA [26]. Each of these algorithms requires tuning of some hyperparameter for unconstrained problems with unknown L_{\max} (usually a scale-factor on a learning rate). In contrast, our RESCALEDEXP requires no such tuning.

We evaluate each algorithm with the average loss after one pass through the data, computing a prediction, an error, and an update to model parameters for each example in the dataset. Note that this is not the same as a cross-validated error, but is closer to the notion of regret addressed in our theorems. We plot this average loss versus hyperparameter setting for each dataset in Figures 1 and 2. These data bear out the effectiveness of RESCALEDEXP: while it is not unilaterally the highest performer on all datasets, it shows remarkable robustness across datasets with zero manual tuning.

4.2 Convolutional Neural Networks

We also evaluated RESCALEDEXP on two convolutional neural network models. These models have demonstrated remarkable success in computer vision tasks and are becoming increasingly more popular in a variety of areas, but can require significant hyperparameter tuning to train. We consider the MNIST [18] and CIFAR-10 [27] image classification tasks.

Our MNIST architecture consisted of two consecutive 5×5 convolution and 2×2 max-pooling layers followed by a 512-neuron fully-connected layer. Our CIFAR-10 architecture was two consecutive 5×5 convolution and 3×3 max-pooling layers followed by a 384-neuron fully-connected layer and a 192-neuron fully-connected layer.




Figure 1: Average loss vs hyperparameter setting for each algorithm across each dataset. RESCALED-EXP has no hyperparameters and so is represented by a flat yellow line. Many of the other algorithms display large sensitivity to hyperparameter setting.

These models are highly non-convex, so that none of our theoretical analysis applies. Our use of RESCALED-EXP is motivated by the fact that in practice convex methods are used to train these models. We found that RESCALED-EXP can match the performance of other popular algorithms (see Figure 3).

In order to achieve this performance, we made a slight modification to RESCALED-EXP: when we update L_t , instead of resetting w_t to zero, we re-center the algorithm about the previous prediction point. We provide no theoretical justification for this modification, but only note that it makes intuitive sense in stochastic optimization problems, where one can reasonably expect that the previous prediction vector is closer to the optimal value than zero.

5 Conclusions

We have presented RESCALED-EXP, an Online Convex Optimization algorithm that achieves regret $\tilde{O}(\|u\| \log(\|u\|) L_{\max} \sqrt{T} + \exp(8 \max_t \|g_t\|^2 / L(t)^2))$ where $L_{\max} = \max_t \|g_t\|$ is unknown in advance. Since RESCALED-EXP does not use any prior-knowledge about the losses or comparison vector u , it is hyperparameter free and so does not require any tuning of learning rates. We also prove a lower-bound showing that any algorithm that addresses the unknown- L_{\max} scenario must suffer an exponential penalty in the regret. We compare RESCALED-EXP to prior optimization algorithms empirically and show that it matches their performance.

While our lower-bound matches our regret bound for RESCALED-EXP in terms of T , clearly there is much work to be done. For example, when RESCALED-EXP is run on the adversarial loss sequence presented in Theorem 1, its regret matches the lower-bound, suggesting that the optimality gap could be improved with superior analysis. We also hope that our lower-bound inspires work in algorithms that adapt to non-adversarial properties of the losses to avoid the exponential penalty.

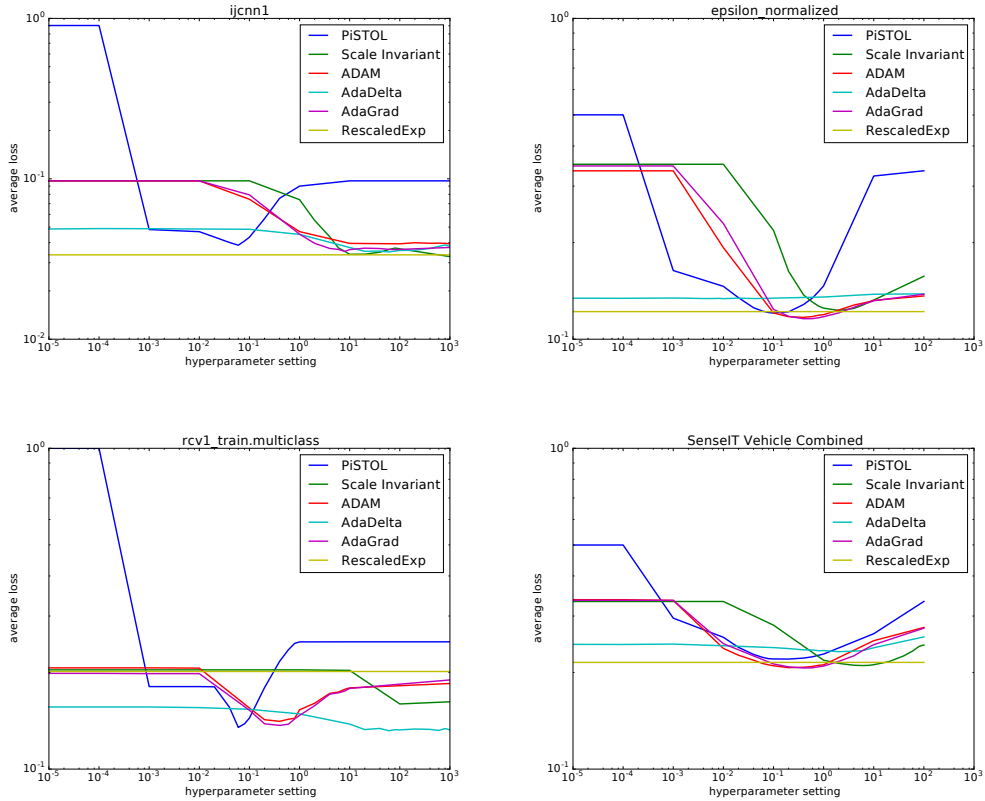


Figure 2: Average loss vs hyperparameter setting, continued from Figure 1.




Figure 3: We compare RESCALED EXP to ADAM, ADAGRAD, and stochastic gradient descent (SGD), with learning-rate hyperparameter optimization for the latter three algorithms. All algorithms achieve a final validation accuracy of 99% on MNIST and 84%, 84%, 83% and 85% respectively on CIFAR-10 (after 40000 iterations).

References

- [1] Martin 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.
- [2] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- [3] Nick Littlestone. From on-line to batch learning. In *Proceedings of the second annual workshop on Computational learning theory*, pages 269–284, 2014.

- [4] Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of on-line learning algorithms. *Information Theory, IEEE Transactions on*, 50(9):2050–2057, 2004.
- [5] J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. In *Conference on Learning Theory (COLT)*, 2010.
- [6] H. Brendan McMahan and Matthew Streeter. Adaptive bound optimization for online convex optimization. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, 2010.
- [7] Brendan McMahan and Matthew Streeter. No-regret algorithms for unconstrained online convex optimization. In *Advances in neural information processing systems*, pages 2402–2410, 2012.
- [8] Francesco Orabona. Dimension-free exponentiated gradient. In *Advances in Neural Information Processing Systems*, pages 1806–1814, 2013.
- [9] Brendan McMahan and Jacob Abernethy. Minimax optimal algorithms for unconstrained linear optimization. In *Advances in Neural Information Processing Systems*, pages 2724–2732, 2013.
- [10] Jacob Abernethy, Peter L Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. In *Proceedings of the nineteenth annual conference on computational learning theory*, 2008.
- [11] Francesco Orabona and Dávid Pál. Scale-free online learning. *arXiv preprint arXiv:1601.01974*, 2016.
- [12] Francesco Orabona and Dávid Pál. Open problem: Parameter-free and scale-free online algorithms. In *Conference on Learning Theory*, 2016.
- [13] S. Shalev-Shwartz. *Online Learning: Theory, Algorithms, and Applications*. PhD thesis, The Hebrew University of Jerusalem, 2007.
- [14] Thomas Hofmann, Bernhard Schölkopf, and Alexander J Smola. Kernel methods in machine learning. *The annals of statistics*, pages 1171–1220, 2008.
- [15] Chih-Chung Chang and Chih-Jen Lin. Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- [16] Isabelle Guyon, Steve Gunn, Asa Ben-Hur, and Gideon Dror. Result analysis of the nips 2003 feature selection challenge. In *Advances in Neural Information Processing Systems*, pages 545–552, 2004.
- [17] Chih-chung Chang and Chih-Jen Lin. Ijcnv 2001 challenge: Generalization ability and text decoding. In *In Proceedings of IJCNN. IEEE*. Citeseer, 2001.
- [18] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [19] David D Lewis, Yiming Yang, Tony G Rose, and Fan Li. Rcv1: A new benchmark collection for text categorization research. *The Journal of Machine Learning Research*, 5:361–397, 2004.
- [20] Marco F Duarte and Yu Hen Hu. Vehicle classification in distributed sensor networks. *Journal of Parallel and Distributed Computing*, 64(7):826–838, 2004.
- [21] M. Lichman. UCI machine learning repository, 2013.
- [22] Shimon Kogan, Dimitry Levin, Bryan R Routledge, Jacob S Sagi, and Noah A Smith. Predicting risk from financial reports with regression. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 272–280. Association for Computational Linguistics, 2009.
- [23] Francesco Orabona, Koby Crammer, and Nicolo Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435, 2014.
- [24] Francesco Orabona. Simultaneous model selection and optimization through parameter-free stochastic learning. In *Advances in Neural Information Processing Systems*, pages 1116–1124, 2014.
- [25] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [26] Matthew D Zeiler. Adadelta: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [27] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images, 2009.
- [28] H. Brendan McMahan. A survey of algorithms and analysis for adaptive online learning. *arXiv preprint arXiv:1403.3465*, 2014.

A Follow-the-Regularized-Leader (FTRL) Regret

Recall that the FTRL algorithm uses the strategy $w_{t+1} = \operatorname{argmin} \psi_t(w) + \sum_{t'=1}^t \ell_{t'}(w)$, where the functions ψ_t are called regularizers.

Theorem 4. *FTRL with regularizers ψ_t and $\psi_0(w_1) = 0$ obtains regret:*

$$R_t(u) \leq \psi_T(u) + \sum_{t=1}^T \psi_{t-1}(w_{t+1}) - \psi_t(w_{t+1}) + \ell_t(w_t) - \ell_t(w_{t+1}) \quad (3)$$

Further, if the losses are linear $\ell_t(w) = g_t \cdot w$ and $\psi_t(w) = \frac{1}{\eta_t} \psi(w)$ for some values η_t and fixed function ψ , then the regret is

$$R_t(u) \leq \frac{1}{\eta_T} \psi(u) + \sum_{t=1}^T \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1}) \quad (4)$$

Proof. The first part follows from some algebraic manipulations:

$$\begin{aligned} \sum_{t=1}^T \ell_t(u) + \psi_T(u) &\geq \psi_T(w_{T+1}) + \sum_{t=1}^T \ell_t(w_{T+1}) \\ - \sum_{t=1}^T \ell_t(u) &\leq \psi_T(u) - \psi_T(w_{T+1}) - \sum_{t=1}^T \ell_t(w_{T+1}) \\ R_T(u) &= \sum_{t=1}^T \ell_t(w_t) - \sum_{t=1}^T \ell_t(u) \\ &\leq \psi_T(u) - \psi_T(w_{T+1}) + \sum_{t=1}^T \ell_t(w_t) - \ell_t(w_{T+1}) \\ &= \psi_T(u) - \psi_T(w_{T+1}) + \ell_T(w_T) - \ell_T(w_{T+1}) + R_{T-1}(w_{T+1}) \\ &\leq \psi_T(u) - \psi_T(w_{T+1}) + \ell_T(w_T) - \ell_T(w_{T+1}) + \\ &\quad + \sum_{t=1}^{T-1} \psi_t(w_{t+2}) - \psi_t(w_{t+1}) + \ell_t(w_t) - \ell_t(w_{t+1}) \\ &= \psi_T(u) + \ell_1(w_1) - \ell_1(w_2) - \psi_1(w_2) \\ &\quad + \sum_{t=2}^T \psi_{t-1}(w_{t+1}) - \psi_t(w_{t+1}) + \ell_t(w_t) - \ell_t(w_{t+1}) \\ &= \psi_T(u) + \sum_{t=1}^T \psi_{t-1}(w_{t+1}) - \psi_t(w_{t+1}) + \ell_t(w_t) - \ell_t(w_{t+1}) \end{aligned}$$

where we're assuming $\psi_0(w_1) = 0$ in the last step.

Now let's specialize to the case of linear losses $\ell_t(w) = g_t \cdot w$ and regularizers of the form $\psi_t(w) = \frac{1}{\eta_t} \psi(w)$ for some fixed regularizer ψ and varying scalings η_t . Plugging this into the previous bound gives:

$$R_t(u) \leq \frac{1}{\eta_T} \psi(u) + \sum_{t=1}^T \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1})$$

□

While this formulation of the regret of FTRL is sufficient for our needs, our analysis is not tight. We refer the reader to [28] for a stronger FTRL bound that can improve constants in some analyses.

B Proof of Lemma 3

We start off by computing the FTRL updates with regularizers $\psi(w)/\eta_t$:

$$\nabla \psi(w) = \log(\|w\| + 1) \frac{w}{\|w\|}$$

so that

$$\begin{aligned} w_{T+1} &= \operatorname{argmin} \frac{1}{\eta_T} \psi(w) + \sum_{t=1}^T g_t \cdot w \\ &= -\frac{g_{1:T}}{\|g_{1:T}\|} (\exp(\eta_T \|g_{1:T}\|) - 1) \end{aligned}$$

Our goal will be to show that the terms $\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t}\right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1})$ in the sum in (4) are negative. In particular, note that sequence of η_t is non-increasing so that $\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t}\right) \psi(w_{t+1}) \leq 0$ for all t . Thus our strategy will be to bound $g_t \cdot (w_t - w_{t+1})$.

B.1 Reduction to one dimension

In order to bound $\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t}\right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1})$, we first show that it suffices to consider the case when g_t and $g_{1:t-1}$ are co-linear.

Theorem 5. *Let W be a separable inner-product space and suppose (with mild abuse of notation) every loss function $\ell_t : W \rightarrow \mathbb{R}$ has some subgradient $g_t \in W^*$ such that $g_t w = \langle g_t, w \rangle$ for some $g_t \in W$. Suppose we run an FTRL algorithm with regularizers $\frac{1}{\eta_t} \psi(\|w\|)$ on loss functions ℓ_t such that $w_{t+1} = \frac{g_{1:t}}{\|g_{1:t}\|} f(\eta_t \|g_{1:t}\|)$ for some function f for all t where $\eta_t = \frac{c}{\sqrt{M_t + \|g_{1:t}\|_1^2}}$ for some constant c . Then for any g_t with $\|g_t\| = L$, both $(\eta_{t-1}^{-1} - \eta_t^{-1}) \psi(\|w_{t+1}\|) + g_t(w_t - w_{t+1})$ and $g_t(w_t - w_{t+1})$ are maximized when g_t is a scalar multiple of $g_{1:t-1}$.*

Proof. The proof is an application of Lagrange multipliers. Our Lagrangian for $(\eta_{t-1}^{-1} - \eta_t^{-1}) \psi(\|w_{t+1}\|) + g_t(w_t - w_{t+1})$ is

$$\begin{aligned} \mathcal{L} &= (\eta_{t-1}^{-1} - \eta_t^{-1}) \psi(\|w_{t+1}\|) + g_t(w_t - w_{t+1}) + \lambda \|g_t\|^2 / 2 \\ &= (\eta_{t-1}^{-1} - \eta_t^{-1}) \psi(f(\eta_t \|g_{1:t}\|)) + g_t \left(w_t - \frac{g_{1:t}}{\|g_{1:t}\|} f(\eta_t \|g_{1:t}\|) \right) + \lambda \frac{\|g_t\|^2}{2} \end{aligned}$$

Fix a countable orthonormal basis of W . For a vector $v \in W$ we let v_i be the projection of v along the i th basis vector of our countable orthonormal basis. We denote the action of $\nabla \mathcal{L}$ on the i th basis vector by $\nabla \mathcal{L}_i$.

Then we have

$$\begin{aligned} \nabla \mathcal{L}_i &= \lambda g_{t,i} + w_{t,i} - w_{t+1,i} - \frac{g_{t,i}}{\|g_{1:t}\|} f(\eta_t \|g_{1:t}\|) \\ &\quad + \sum_j \frac{g_{t,j} (g_{1:t})_j}{\|g_{1:t}\|^3} (g_{1:t})_i f(\eta_t \|g_{1:t}\|) \\ &\quad - \sum_j \frac{(g_{1:t})_j g_{t,j}}{\|g_{1:t}\|} f'(\eta_t \|g_{1:t}\|) \left[\frac{(g_{1:t})_i \eta_t}{\|g_{1:t}\|} - \frac{\|g_{1:t}\| c \left(\frac{\partial M_t}{\partial g_{t,i}} + 2g_{t,i} \right)}{2(M_t + \|g_{1:t}\|_1^2)^{3/2}} \right] \\ &\quad + (\eta_{t-1}^{-1} - \eta_t^{-1}) \psi'(f(\eta_t \|g_{1:t}\|)) f'(\eta_t \|g_{1:t}\|) \left[\frac{(g_{1:t})_i \eta_t}{\|g_{1:t}\|} - \frac{\|g_{1:t}\| c \left(\frac{\partial M_t}{\partial g_{t,i}} + 2g_{t,i} \right)}{2(M_t + \|g_{1:t}\|_1^2)^{3/2}} \right] \\ &\quad - \psi(f(\eta_t \|g_{1:t}\|)) \frac{\frac{\partial M_t}{\partial g_{t,i}} + 2g_{t,i}}{2c \sqrt{M_t + \|g_{1:t}\|_1^2}} \\ &= \lambda g_{t,i} + w_{t,i} - w_{t+1,i} + A g_{t,i} + B (g_{1:t-1})_i + C \frac{\partial M_t}{\partial g_{t,i}} \end{aligned}$$

where A , B and C do not depend on i . Since $w_{t,i}$ and $w_{t+1,i}$ are scalar multiples of $g_{1:t-1}$ and $g_{1:t}$ respectively, we can reassign the variables A and B to write

$$\nabla \mathcal{L}_i = A g_{t,i} + B (g_{1:t-1})_i + C \frac{\partial M_t}{\partial g_{t,i}}$$

Now we compute

$$\begin{aligned}\frac{\partial M_t}{\partial g_{t,i}} &= \frac{\partial \max(M_{t-1}, \|g_{1:t}\|/p - \|g_{1:t}\|_1^2)}{\partial g_{t,i}} \\ &= \begin{cases} 0 & : M_t = M_{t-1} \\ \frac{(g_{1:t})_i}{p\|g_{1:t}\|} - 2g_{t,i} & : M_t \neq M_{t-1} \end{cases}\end{aligned}$$

Thus after again reassigning the variables A and B we have

$$\nabla \mathcal{L}_i = Ag_{t,i} + B(g_{1:t-1})_i$$

Therefore we can only have $\nabla \mathcal{L} = 0$ if g_t is a scalar multiple of $g_{1:t-1}$ as desired.

For $g_t(w_t - w_{t+1})$, we apply exactly the same argument. The Lagrangian is

$$\begin{aligned}\mathcal{L} &= g_t(w_t - w_{t+1}) + \lambda \|g_t\|^2/2 \\ &= g_t \left(w_t - \frac{g_{1:t}}{\|g_{1:t}\|} f(\eta_t \|g_{1:t}\|) \right) + \lambda \frac{\|g_t\|^2}{2}\end{aligned}$$

and differentiating we have

$$\begin{aligned}\nabla \mathcal{L}_i &= \lambda g_{t,i} + w_{t,i} - w_{t+1,i} - \frac{g_{t,i}}{\|g_{1:t}\|} f(\eta_t \|g_{1:t}\|) \\ &\quad + \sum_j \frac{g_{t,j}(g_{1:t})_j}{\|g_{1:t}\|^3} (g_{1:t})_i f(\eta_t \|g_{1:t}\|) \\ &\quad - \sum_j \frac{(g_{1:t})_j g_{t,j}}{\|g_{1:t}\|} f'(\eta_t \|g_{1:t}\|) \left[\frac{(g_{1:t})_i \eta_t}{\|g_{1:t}\|} - \frac{\|g_{1:t}\| c \left(\frac{\partial M_t}{\partial g_{t,i}} + 2g_{t,i} \right)}{2(M_t + \|g_{1:t}\|_1^2)^{3/2}} \right] \\ &= \lambda g_{t,i} + w_{t,i} - w_{t+1,i} + Ag_{t,i} + B(g_{1:t-1})_i + C \frac{\partial M_t}{\partial g_{t,i}} \\ &= Ag_{t,i} + B(g_{1:t-1})_i\end{aligned}$$

so that again we are done. □

We make the following intuitive definition:

Definition 6. For any vector $v \in W$, define $\text{sign}(v) = \frac{v}{\|v\|}$.

In the next section, we prove bounds on the quantity $(\eta_{t-1}^{-1} - \eta_t^{-1})\psi(\|w_{t+1}\|) + g_t(w_t - w_{t+1})$. By Theorem 5 this quantity is maximized when $\text{sign}(g_t) = \pm \text{sign}(g_{1:t-1})$ and so we consider only this case.

B.2 One dimensional FTRL

In this section we analyze the regret of our FTRL algorithm with the end-goal of proving Lemma 3. We make heavy use of Theorem 5 to allow us to consider only the case $\text{sign}(g_t) = \pm \text{sign}(g_{1:t-1})$. In this setting we may identify the 1-dimensional space spanned by g_t and $g_{1:t-1}$ with \mathbb{R} . Thus whenever we are operating under the assumption $\text{sign}(g_t) = \text{sign}(g_{1:t-1})$ we will use $|\cdot|$ in place of $\|\cdot\|$ and occasionally assume $g_{1:t-1} > 0$ as this holds WLOG. We feel that this notation and assumption aids intuition in visualizing the following results.

Lemma 7. Suppose $\text{sign}(g_t) = \text{sign}(g_{1:t-1})$. Then

$$|\eta_{t-1}\|g_{1:t-1}\| - \eta_t\|g_{1:t}\|| \leq \eta_t\|g_t\| \quad (5)$$

Suppose instead that $\text{sign}(g_t) = -\text{sign}(g_{1:t-1})$ and also $\|g_t\| \leq L$. Then we still have:

$$|\eta_{t-1}\|g_{1:t-1}\| - \eta_t\|g_{1:t}\|| \leq \left(1 + \frac{pL}{2}\right) \eta_t\|g_t\| \quad (6)$$

Proof. First, suppose $\text{sign}(g_t) = \text{sign}(g_{1:t-1})$. Then $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$. WLOG, assume $g_{1:t-1} > 0$. Notice that $\eta_t g_{1:t}$ is an increasing function of g_t for $g_t > 0$ because $\eta_t g_{1:t}$ is proportional to either $g_{1:t}$ or $\sqrt{g_{1:t}}$ depending on whether $M_t = M_{t-1}$ or not. Then since $\eta_t < \eta_{t-1}$ we have

$$\begin{aligned}|\eta_{t-1}g_{1:t-1} - \eta_t g_{1:t}| &= \eta_t g_{1:t} - \eta_{t-1}g_{1:t-1} \\ &\leq \eta_t g_{1:t} - \eta_t g_{1:t-1} \\ &= \eta_t |g_t|\end{aligned}$$

so that (5) holds.

Now suppose $\text{sign}(g_t) = -\text{sign}(g_{1:t-1})$ and $\|g_t\| \leq L$. We consider two cases.

Case 1: $\eta_t |g_{1:t}| \geq \eta_{t-1} |g_{1:t-1}|$:

Since $\eta_{t-1} \geq \eta_t$, we have

$$\begin{aligned}\eta_t |g_{1:t}| &\geq \eta_{t-1} |g_{1:t-1}| \\ \eta_t |g_{1:t}| &\geq \eta_t |g_{1:t-1}| \\ |g_{1:t}| &\geq |g_{1:t-1}| \\ |g_t| &\geq |g_{1:t}|\end{aligned}$$

where the last line follows since $\text{sign}(g_{1:t-1}) = -\text{sign}(g_t)$. Therefore:

$$|\eta_{t-1} |g_{1:t-1}| - \eta_t |g_{1:t}| \leq \eta_t |g_{1:t}| \leq \eta_t |g_t|$$

so that we are done.

Case 2: $\eta_t |g_{1:t}| \leq \eta_{t-1} |g_{1:t-1}|$:

When $g_t < -g_{1:t-1}$ and $\eta_t |g_{1:t}| \leq \eta_{t-1} |g_{1:t-1}|$, $|\eta_{t-1} |g_{1:t-1}| - \eta_t |g_{1:t}|$ is a decreasing function of $|g_t|$ because $\eta_t |g_{1:t}|$ is an increasing function of $|g_t|$ for $g_t < -g_{1:t-1}$. Therefore it suffices to consider the case $g_t \geq -g_{1:t-1}$, so that $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$ and $|g_{1:t}| \leq |g_{1:t-1}|$:

Since $|g_{1:t}| \leq |g_{1:t-1}|$, we have $M_t = M_{t-1}$ so that we can write:

$$\begin{aligned}\eta_{t-1} g_{1:t-1} - \eta_t g_{1:t} &= -g_t \eta_t + g_{1:t-1} (\eta_{t-1} - \eta_t) \\ &= |g_t| \eta_t + g_{1:t-1} \left(\frac{1}{k\sqrt{2}\sqrt{M_{t-1} + \|g\|_{1:t-1}^2}} - \frac{1}{k\sqrt{2}\sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2}} \right) \\ &= |g_t| \eta_t + \frac{g_{1:t-1}}{k\sqrt{2}} \left(\frac{1}{\sqrt{M_{t-1} + \|g\|_{1:t-1}^2}} - \frac{1}{\sqrt{M_{t-1} + \|g\|_{1:t-1}^2 + g_t^2}} \right) \\ &\leq |g_t| \eta_t + \frac{g_{1:t-1}}{k\sqrt{2}\sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2}} \left(\sqrt{M_{t-1} + \|g\|_{1:t-1}^2 + g_t^2} - 1 \right) \\ &\leq |g_t| \eta_t + g_{1:t-1} \eta_t \left(1 + \frac{g_t^2}{2(M_{t-1} + \|g\|_{1:t-1}^2)} - 1 \right) \\ &\leq |g_t| \eta_t + \eta_t \frac{g_{1:t-1} g_t^2}{2(M_{t-1} + \|g\|_{1:t-1}^2)} \\ &\leq |g_t| \eta_t \left(1 + \frac{pL}{2} \right)\end{aligned}$$

we have used the identity $\sqrt{X + g_t^2} \leq \sqrt{X} + \frac{g_t^2}{2\sqrt{X}}$ between lines 4 and 5, and the last line follows because $|g_t| \leq L$ and $M_{t-1} + \|g\|_{1:t-1}^2 \geq |g_{1:t-1}|/p$. \square

Lemma 8. *If*

$$\|w_T\| \geq \exp\left(\frac{\sqrt{pB}}{k\sqrt{2}}\right) - 1$$

then

$$\|g_{1:T-1}\| \geq B$$

Proof. First note that by definition of M_{T-1} and η_{T-1} , $\eta_{T-1} \|g_{1:T-1}\| \leq \frac{\sqrt{p\|g_{1:T-1}\|}}{k\sqrt{2}}$. The proof now follows from some algebra:

$$\begin{aligned}\exp\left(\frac{\sqrt{pB}}{k\sqrt{2}}\right) &\leq \|w_T\| + 1 \\ &= \exp(\eta_{T-1} \|g_{1:T-1}\|) \\ &\leq \exp\left(\frac{\sqrt{p\|g_{1:T-1}\|}}{k\sqrt{2}}\right)\end{aligned}$$

Taking squares of logs and rearranging now gives the desired inequality. \square

We have the following immediate corollary:

Corollary 9. *Suppose $\text{sign}(g_t) = \pm \text{sign}(g_{1:t-1})$, $\|g_t\| \leq L$, and*

$$\|w_t\| \geq \exp\left(\frac{\sqrt{pL}}{k\sqrt{2}}\right) - 1$$

Then $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$.

Now we begin analysis of the sum term in (4).

Lemma 10. *Suppose $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$ and $|g_t| \leq L$. Then*

$$|w_t - w_{t+1}| \leq |g_t| \eta_t (|w_{t+1}| + 1) \left(1 + \frac{pL}{2}\right) \exp\left[g_t \eta_t \left(1 + \frac{pL}{2}\right)\right]$$

Proof. Since $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$, we have:

$$\begin{aligned} |w_t - w_{t+1}| &= |\text{sign}(g_{1:t-1}) [\exp(\eta_{t-1}|g_{1:t-1}|) - 1] - [\text{sign}(g_{1:t}) \exp(\eta_t|g_{1:t}|) - 1]| \\ &= |\exp(\eta_{t-1}|g_{1:t-1}|) - \exp(\eta_t|g_{1:t}|)| \\ &= (|w_{t+1}| + 1) |\exp(\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|) - 1| \end{aligned}$$

where the last line uses the definition of w_{t+1} to observe that $|w_{t+1}| + 1 = \exp(\eta_t|g_{1:t}|)$. Now we consider two cases: either $\eta_{t-1}|g_{1:t-1}| < \eta_t|g_{1:t}|$ or not.

Case 1: $\eta_{t-1}|g_{1:t-1}| < \eta_t|g_{1:t}|$:

By convexity of exp, we have

$$\begin{aligned} |w_t - w_{t+1}| &\leq (|w_{t+1}| + 1) |\exp(\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|) - 1| \\ &\leq (|w_{t+1}| + 1) |\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|| \\ &\leq (|w_{t+1}| + 1) \left(1 + \frac{pL}{2}\right) \eta_t |g_t| \end{aligned}$$

so that the lemma holds.

Case 2: $\eta_{t-1}|g_{1:t-1}| \geq \eta_t|g_{1:t}|$:

Again by convexity of exp we have

$$\begin{aligned} |w_t - w_{t+1}| &\leq (|w_{t+1}| + 1) |\exp(\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|) - 1| \\ &\leq (|w_{t+1}| + 1) |\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|| \exp(\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|) \\ &\leq (|w_{t+1}| + 1) \left(1 + \frac{pL}{2}\right) \exp\left[\eta_t|g_t| \left(1 + \frac{pL}{2}\right)\right] \eta_t |g_t| \end{aligned}$$

so that the lemma still holds. □

The next lemma is the main workhorse of our regret bounds:

Lemma 11. *Suppose $\|g_t\| \leq L$ and either of the following holds:*

1. $p \leq \frac{2}{L}$, $k = \sqrt{2}$, and $\|w_t\| \geq 15$.
2. $k = \sqrt{2}$, $pL \geq 1$, and $\|w_t\| \geq 4 \exp(p^2 L^2)$.

Then

$$\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t}\right) \psi(w_{t+1}) + g_t (w_t - w_{t+1}) \leq 0 \tag{7}$$

Further, inequality (7) holds for any k and sufficiently large L if $\|w_t\| \geq \exp((pL)^2)$.

Proof. By Theorem 5 it suffices to consider the case $\text{sign}(g_t) = \pm \text{sign}(g_{1:t-1})$, so that we may adopt our identification with \mathbb{R} and use of $|\cdot|$ throughout this proof.

For $p \leq \frac{2}{L}$, $k = \sqrt{2}$ we have $15 > \exp\left(\frac{\sqrt{pL}}{k\sqrt{2}}\right) - 1$ and for sufficiently large L , $\exp((pL)^2) > \exp\left(\frac{\sqrt{pL}}{k\sqrt{2}}\right) - 1$. Therefore in all cases $|w_t| \geq \exp\left(\frac{\sqrt{pL}}{k\sqrt{2}}\right) - 1$ so that by Corollary 9 and Lemma 10 we have

$$g_t \cdot (w_t - w_{t+1}) \leq \eta_t g_t^2 (|w_{t+1}| + 1) \left(1 + \frac{pL}{2}\right) \exp\left[\eta_t g_t \left(1 + \frac{pL}{2}\right)\right] \tag{8}$$

First, we prove that (7) is guaranteed if the following holds:

$$|w_{t+1}| + 1 \geq \exp \left[\frac{1 + \frac{pL}{2}}{k^2} \exp \left(\eta_t g_t \left(1 + \frac{pL}{2} \right) \right) + 1 \right] \quad (9)$$

The previous line (9) is equivalent to:

$$k^2 (\log(|w_{t+1}| + 1) - 1) \geq \left(1 + \frac{pL}{2} \right) \exp \left(\eta_t g_t \left(1 + \frac{pL}{2} \right) \right) \quad (10)$$

Notice that $\psi(w_{t+1}) = (|w_{t+1}| + 1)(\log(|w_{t+1}| + 1) - 1) + 1 \geq (|w_{t+1}| + 1)(\log(|w_{t+1}| + 1) - 1)$. Then multiplying (10) by $\eta_t |g_t|$ we have

$$(|w_{t+1}| + 1) \left(1 + \frac{pL}{2} \right) \exp \left[\eta_t |g_t| \left(1 + \frac{pL}{2} \right) \right] \eta_t |g_t| \leq k^2 \eta_t |g_t| \psi(w_{t+1}) \quad (11)$$

Combining (8) and (11), we see that (9) implies

$$g_t \cdot (w_t - w_{t+1}) \leq k^2 \eta_t g_t^2 \psi(w_{t+1})$$

Now we bound $\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1})$:

$$\begin{aligned} \frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} &= k\sqrt{2} \left(\sqrt{M_{t-1} + \|g\|_{1:t-1}^2} - \sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2} \right) \\ &\leq k\sqrt{2} \left(\left[\sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2} - \frac{g_t^2 + M_t - M_{t-1}}{2\sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2}} \right] - \sqrt{M_t + \|g\|_{1:t-1}^2 + g_t^2} \right) \\ &\leq -k\sqrt{2} \frac{g_t^2}{2\sqrt{M_t + \|g\|_{1:t}^2}} \\ &= -k^2 \eta_t g_t^2 \end{aligned}$$

Thus when (9) holds we have

$$\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t (w_t - w_{t+1}) \leq -k^2 \eta_t g_t^2 \psi(w_{t+1}) + k^2 \eta_t^2 g_t^2 \psi(w_{t+1}) \leq 0$$

Therefore our objective is to show that our conditions on w_t imply the condition (9) on w_{t+1} .

First, we bound $\eta_t g_t$ in terms of $|w_t|$. Notice that

$$\begin{aligned} |w_t| + 1 &= \exp \left(\frac{|g_{1:t-1}|}{k\sqrt{2}\sqrt{M_{t-1} + \|g\|_{1:t-1}^2}} \right) \\ &\leq \exp \left(\frac{\sqrt{p}\sqrt{|g_{1:t-1}|}}{k\sqrt{2}} \right) \\ \frac{2k^2 \log^2(|w_t| + 1)}{p} &\leq |g_{1:t-1}| \end{aligned}$$

Using this we have:

$$\begin{aligned} \eta_t g_t &= \frac{g_t}{k\sqrt{2}\sqrt{M_t + \|g\|_{1:t}^2}} \\ &\leq \frac{g_t}{k\sqrt{2}\sqrt{M_{t-1} + \|g\|_{1:t-1}^2 + g_t^2}} \\ &\leq \frac{g_t \sqrt{p}}{k\sqrt{2}\sqrt{|g_{1:t-1}|} + pg_t^2} \\ &\leq \frac{L\sqrt{p}}{k\sqrt{2}\sqrt{\frac{2k^2}{p} \log^2(|w_t| + 1) + pL^2}} \end{aligned}$$

so that we can conclude:

$$\eta_t g_t \leq \frac{Lp}{k\sqrt{2}\sqrt{2k^2 \log^2(|w_t| + 1) + p^2 L^2}} \quad (12)$$

Further, by Lemma 7 we have

$$\begin{aligned} \frac{|w_t| + 1}{|w_{t+1}| + 1} &= \exp(\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}|) \\ &\leq \exp\left[\eta_t g_t \left(1 + \frac{pL}{2}\right)\right] \end{aligned}$$

Therefore we have

$$|w_{t+1}| + 1 \geq (|w_t| + 1) \exp\left[-\eta_t g_t \left(1 + \frac{pL}{2}\right)\right] \quad (13)$$

From (13), we see that (9) is guaranteed if we have

$$|w_t| + 1 \geq \exp\left[\eta_t g_t \left(1 + \frac{pL}{2}\right)\right] \exp\left[\frac{1 + \frac{pL}{2}}{k^2} \exp\left(\eta_t g_t \left(1 + \frac{pL}{2}\right)\right) + 1\right] \quad (14)$$

If we use our expression (12) in (14), and assume $|w_t| \geq \exp(L^2)$, we see that there exists some constant C depending on p and k such that the RHS of (14) is $O(\exp(L))$ and so (14) holds for sufficiently large L .

For $p = 2/L$, $k = \sqrt{2}$, and $w_t \geq 15$ we can verify (14) numerically by plugging in the bound (12).

For the case $k = \sqrt{2}$, $|w_t| \geq 4 \exp(p^2 L^2)$, we notice that by using (12), we can write (14) entirely in terms of pL . Graphing both sides numerically as functions of pL then allows us to verify the condition. \square

We have one final lemma we need before we can start stating some real regret bounds. This lemma can be viewed as observing that $\psi(w)$ is roughly $\frac{1}{D}$ strongly-convex for $|w|$ not much bigger than D .

Lemma 12. *Suppose $p \leq 2/L$, $k = \sqrt{2}$, $\|w_t\| \leq D$ and $\|g_t\| \leq L$ Then $g_t(w_t - w_{t+1}) \leq 6(\max(D + 1, \exp(1/2)))g_t^2 \eta_t$.*

Proof. By Theorem 5 it suffices to consider $\text{sign}(g_t) = \pm \text{sign}(g_{1:t-1})$.

We show that $|w_t - w_{t+1}| \leq 6(\max(D + 1, \exp(1/2)))|g_t|\eta_t$ so that the result follows by multiplying by $|g_t|$.

From Lemma 7, we have $|\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}| \leq \eta_t|g_t| \left(1 + \frac{pL}{2}\right) \leq 2\eta_t|g_t|$. Further, note that $\eta_t|g_t| \leq \frac{1}{k\sqrt{2}} = \frac{1}{2}$. We consider two cases, either $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$ or not.

Case 1: $\text{sign}(g_{1:t}) = \text{sign}(g_{1:t-1})$:

$$\begin{aligned} |w_t - w_{t+1}| &= |\exp(\eta_{t-1}|g_{1:t-1}|) - \exp(\eta_t|g_{1:t}|)| \\ &= (|w_t| + 1) |\exp(\eta_t|g_{1:t}| - \eta_{t-1}|g_{1:t-1}|) - 1| \\ &\leq 2(D + 1)\eta_t|g_t| \exp(2\eta_t|g_t|) \\ &\leq 2(D + 1)\eta_t|g_t| \exp\left(\frac{2}{k\sqrt{2}}\right) \\ &\leq 6(D + 1)\eta_t|g_t| \end{aligned}$$

Case 2: $\text{sign}(g_{1:t}) \neq \text{sign}(g_{1:t-1})$: In this case, we must have $|g_{1:t}| \leq |g_t|$. Let $X = \max(\eta_t|g_{1:t}|, \eta_{t-1}|g_{1:t-1}|)$. Then by triangle inequality we have

$$\begin{aligned} |w_t - w_{t+1}| &\leq 2 \max(|w_t|, |w_{t+1}|) \\ &\leq 2(\exp(X) - 1) \\ &\leq 2X \exp(X) \\ &\leq 2(\max(|w_t|, |w_{t+1}|) + 1)X \end{aligned}$$

Since $|\eta_{t-1}|g_{1:t-1}| - \eta_t|g_{1:t}| \leq 2\eta_t g_t$, we have $X \leq 2\eta_t g_t + \eta_t|g_{1:t}| \leq 3\eta_t|g_t|$ so that we have

$$|w_t - w_{t+1}| \leq 6(\max(|w_t|, |w_{t+1}|) + 1)\eta_t|g_t|$$

Finally, we have $|w_{t+1}| + 1 = \exp(\eta_t |g_{1:t}|) \leq \exp(\eta_t |g_t|) \leq \exp(1/2)$, so that

$$\begin{aligned} |w_t - w_{t+1}| &\leq 6\eta_t |g_t| (\max(|w_t|, |w_{t+1}|) + 1) \\ &\leq 6 \max(D + 1, \exp(1/2)) \eta_t |g_t| \end{aligned}$$

□

Now we are finally in a position to prove Lemma 3, which we re-state below:

Lemma 3. *Set $k = \sqrt{2}$. Suppose $\|g_t\| \leq L$ for $t < T$, $1/L \leq p \leq 2/L$, $g_T \leq L_{\max}$ and $L_{\max} \geq L$. Let $W_{\max} = \max_{t \in [1, T]} \|w_t\|$. Then the regret of FTRL with regularizers $\psi_t(w) = \psi(w)/\eta_t$ is:*

$$\begin{aligned} R_T(u) &\leq \psi(u)/\eta_T + 96\sqrt{M_T + \|g\|_{1:T}^2} + 2L_{\max} \min \left[W_{\max}, 4 \exp \left(4 \frac{L_{\max}^2}{L^2} \right), \exp(\sqrt{T/2}) \right] \\ &\leq (2\psi(u) + 96) \sqrt{\sum_{t=1}^{T-1} L|g_t| + L_{\max}^2} + 8L_{\max} \min \left[\exp \left(\frac{4L_{\max}^2}{L^2} \right), \exp(\sqrt{T/2}) \right] \\ &\leq L_{\max} (2(\|u\| + 1) \log(\|u\| + 1) - \|u\|) + 96\sqrt{T} + 8L_{\max} \min \left[e^{\frac{4L_{\max}^2}{L^2}}, e^{\sqrt{T/2}} \right] \end{aligned}$$

Proof of Lemma 3. We combine Lemma 11 with Lemma 12: if $|w_t| \geq 15$ we have for all $t < T$:

$$\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1}) < 0$$

and if $|w_t| \leq 15$ we have

$$\begin{aligned} \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1}) &\leq g_t \cdot (w_t - w_{t+1}) \\ &\leq 6 \times (15 + 1) \eta_t g_t^2 \\ &= 96 \eta_t g_t^2 \end{aligned}$$

Therefore for all $t < T$ we have $\left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1}) \leq 96 \eta_t g_t^2$.

$$\begin{aligned} R_T(u) &\leq \psi(u)/\eta_T + \sum_{t=1}^T \left(\frac{1}{\eta_{t-1}} - \frac{1}{\eta_t} \right) \psi(w_{t+1}) + g_t \cdot (w_t - w_{t+1}) \\ &\leq \psi(u)/\eta_T + 96 \sum_{t=1}^T \eta_t g_t^2 + \left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) + g_T \cdot (w_T - w_{T+1}) \end{aligned}$$

We have

$$\left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) < 0$$

so that

$$\left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) + g_T \cdot (w_T - w_{T+1}) \leq 2L_{\max} W_{\max}$$

Further, again using Lemma 11 we have

$$\left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) + g_T \cdot (w_T - w_{T+1}) < 0$$

for $|w_T| \geq 4 \exp(p^2 L_{\max}^2)$ since $k = \sqrt{2}$.

Finally, notice that by definition of η_t and L , we must have $|\eta_t g_{1:t}| \leq \frac{\sqrt{p|g_{1:t}|}}{k\sqrt{2}} \leq \sqrt{T/2}$, so that $\|w_t\| \leq \exp(\eta_t |g_{1:t}|) \leq \exp(\sqrt{T/2})$. Thus we have

$$\left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) + g_T \cdot (w_T - w_{T+1}) \leq 2L_{\max} \min(W_{\max}, 4 \exp(4L_{\max}^2/L^2), \exp(\sqrt{2T}))$$

Now we make the following classic argument:

$$\begin{aligned} \sqrt{M_t + \|g\|_{1:t}^2} - \sqrt{M_{t-1} + \|g\|_{1:t-1}^2} &\geq \frac{g_t^2 + M_t - M_{t-1}}{2\sqrt{M_t + \|g\|_{1:t}^2}} \\ &\geq \frac{g_t^2}{2\sqrt{M_t + \|g\|_{1:t}^2}} \\ &= \eta_t g_t^2 \end{aligned}$$

so that we can bound:

$$\begin{aligned} R_T(u) &\leq \psi(u)/\eta_T + 96 \sum_{t=1}^T \eta_t g_t^2 + \left(\frac{1}{\eta_{T-1}} - \frac{1}{\eta_T} \right) \psi(w_{T+1}) + g_T \cdot (w_T - w_{T+1}) \\ &\leq \psi(u)/\eta_T + 96 \sqrt{M_T + \|g\|_{1:T}^2} + 2L_{\max} \min(W_{\max}, 4 \exp(4L_{\max}^2/L^2), \exp(\sqrt{2T})) \end{aligned}$$

To show the remaining two lines of the theorem, we prove by induction that $M_t + \|g\|_{1:t}^2 \leq L \sum_{t'=1}^t |g_{t'}|$ for all $t < T$. The statement is clearly true for $t = 1$. Suppose it holds for some t . Then notice that $|g_{1:t+1}| \leq |g_{t+1}| + |g_{1:t}|$. So we have

$$\begin{aligned} M_{t+1} + \|g\|_{1:t+1}^2 &= \max \left(M_t + \|g\|_{1:t+1}^2, \frac{|g_{1:t+1}|}{p} \right) \\ &\leq \max \left(M_t + \|g\|_{1:t}^2 + L|g_{t+1}|, L|g_{1:t+1}| \right) \\ &\leq L \sum_{t'=1}^{t+1} |g_{t'}| \end{aligned}$$

Finally, we observe that $M_T = \max \left(M_{T-1} + \|g\|_{1:T-1}^2 + g_T^2, \frac{|g_{1:T}|}{p} \right) \leq L_{\max}^2 + L \sum_{t=1}^{T-1} |g_{t'}|$ and the last two lines of the theorem follow immediately. \square

C Additional Experimental Details

C.1 Hyperparameter Optimization

For the linear classification tasks, we optimized hyperparameters in a two-step process. First, we tested every power of 10 from 10^{-5} to 10^2 . Second, if λ was the best hyperparameter setting in step 1, we additionally tested $\beta\lambda$ for $\beta \in \{0.2, 0.4, 0.8, 2.0, 4.0, 6.0, 8.0\}$

For the neural network models, we optimized ADAM and ADAGRAD's learning rates by testing every power of 10 from 10^{-5} to 10^0 . For stochastic gradient descent, we used an exponentially decaying learning rate schedule specified in Tensorflow's (<https://www.tensorflow.org/>) MNIST and CIFAR-10 example code.

C.2 Coordinate-wise updates

We proved all our results in arbitrarily many dimensions, leading to a dimension-independent regret bound. However, it is also possible to achieve dimension-dependent bounds by running an independent version of our algorithm on each coordinate. Formally, for OLO we have

$$R_T(u) = \sum_{t=1}^T g_t(w_t - u) = \sum_{i=1}^d \sum_{t=1}^T g_{t,i}(w_{t,i} - u_i) = \sum_{i=1}^d R_T^1(u_i)$$

where R_T^1 is the regret of a 1-dimensional instance of the algorithm. This reduction can yield substantially better regret bounds when the gradients g_t are known to be sparse (but can be much worse when they are not). We use this coordinate-wise update strategy for our linear classification experiments for RESCALEDEXP. We also considered coordinate-wise updates and non-coordinate wise updates for the other algorithms, taking the best-performing of the two.

For all algorithms in the linear classification experiments, we found that the difference between coordinate-wise and non-coordinate wise updates was not very striking. However, for the neural network experiments we found RESCALEDEXP performed extremely poorly when using coordinate-wise updates, and performed extremely well with non-coordinate wise updates. We hypothesize that this is due to a combination of non-convexity of the model and frequent resets at different times for each coordinate.

ADAGRAD	RESCALEDEXP	ADADELTA	SCALEINVARIANT	ADAM	PISTOL
1.14	1.19	1.21	1.28	1.51	1.53

Table 1: Average normalized loss, using best hyperparameter setting for each algorithm.

C.3 Re-centering RESCALEDEXP

For the non-convex neural network tasks we used a variant of RESCALEDEXP in which we re-center our FTRL algorithm at the beginning of each epoch. Formally, the pseudo-code is provided below:

Algorithm 2 Re-centered RESCALEDEXP

Initialize: $k \leftarrow \sqrt{2}$, $M_0 \leftarrow 0$, $w_1 \leftarrow 0$, $t_* \leftarrow 1$, $w_* \leftarrow 0$
for $t = 1$ **to** T **do**
 Play w_t , receive subgradient $g_t \in \partial \ell_t(w_t)$.
 if $t = 1$ **then**
 $L_1 \leftarrow \|g_1\|$
 $p \leftarrow 1/L_1$
 end if
 $M_t \leftarrow \max(M_{t-1}, \|g_{t_*:t}\|/p - \|g\|_{t_*:t}^2)$.
 $\eta_t \leftarrow \frac{1}{k\sqrt{2(M_t + \|g\|_{t_*:t}^2)}}$
 $w_{t+1} \leftarrow w_* + \operatorname{argmin}_w \left[\frac{\psi(w)}{\eta_t} + g_{t_*:t} w \right] = w_* - \frac{g_{t_*:t}}{\|g_{t_*:t}\|} [\exp(\eta_t \|g_{t_*:t}\|) - 1]$
 if $\|g_t\| > 2L_t$ **then**
 $L_{t+1} \leftarrow \|g_t\|$
 $p \leftarrow 1/L_{t+1}$
 $t_* \leftarrow t + 1$
 $M_t \leftarrow 0$
 $w_{t+1} \leftarrow 0$
 $w_* \leftarrow w_{t-1}$
 else
 $L_{t+1} \leftarrow L_t$
 end if
end for

So long as $\|w_* - u\| \leq \|u\|$, this algorithm maintains the same regret bound as the non-re-centered version of RESCALEDEXP. While it is intuitively reasonable to expect this to occur in a stochastic setting, an adversary can easily subvert this algorithm.

C.4 Aggregating Studies

It is difficult to interpret the results of a study such as our linear classification experiments (see Section 4) in which no particular algorithm is always the “winner” for every dataset. In particular, consider the case of an analyst who wishes to run one of these algorithms on some new dataset, and doesn’t have the either the resources or inclination to implement and tune each algorithm. Which should she choose? We suggest the following heuristic: pick the algorithm with the lowest loss averaged *across datasets*.

This heuristic is problematic because datasets in which all algorithms do very poorly will dominate the cross-dataset average. In order address this issue and compare losses across datasets properly, we compute a *normalized loss* for each algorithm and dataset. The normalized loss for an algorithm on a dataset is given by taking the loss experienced by the algorithm on its best hyperparameter setting on that dataset divided by the lowest loss observed by any algorithm and hyperparameter setting on that dataset. Thus a normalized loss of 1 on a dataset indicates that an algorithm outperformed all other algorithms on the dataset (at least for its best hyperparameter setting). We then average the normalized loss for each algorithm across datasets to obtain the scores for each algorithm (see Table 1).

These data indicate that while ADAGRAD has a slight edge after tuning, RESCALEDEXP and ADADELTA do nearly equivalently well (4% and 6% worse performance, respectively). Therefore we suggest that if our intrepid analyst is willing to perform some hyperparameter tuning, then ADAGRAD may be slightly better, but her choice doesn’t matter too much. On the other hand, using RESCALEDEXP will allow her to skip any tuning step without compromising performance.