
Dimension-Free Exponentiated Gradient

Francesco Orabona

Toyota Technological Institute at Chicago
Chicago, USA

francesco@orabona.com

Abstract

I present a new online learning algorithm that extends the exponentiated gradient framework to infinite dimensional spaces. My analysis shows that the algorithm is implicitly able to estimate the L_2 norm of the unknown competitor, U , achieving a regret bound of the order of $\mathcal{O}(U \log(UT + 1))\sqrt{T}$, instead of the standard $\mathcal{O}((U^2 + 1)\sqrt{T})$, achievable without knowing U . For this analysis, I introduce novel tools for algorithms with time-varying regularizers, through the use of local smoothness. Through a lower bound, I also show that the algorithm is optimal up to $\sqrt{\log(UT)}$ term for linear and Lipschitz losses.

1 Introduction

Online learning provides a scalable and flexible approach for solving a wide range of prediction problems, including classification, regression, ranking, and portfolio management. These algorithms work in rounds, where at each round a new instance is given and the algorithm makes a prediction. After the true label of the instance is revealed, the learning algorithm updates its internal hypothesis. The aim of the classifier is to minimize the cumulative loss it suffers due to its prediction, such as the total number of mistakes.

Popular online algorithms for classification include the standard Perceptron and its many variants, such as kernel Perceptron [6], and p -norm Perceptron [7]. Other online algorithms, with properties different from those of the standard Perceptron, are based on multiplicative (rather than additive) updates, such as Winnow [10] for classification and Exponentiated Gradient (EG) [9] for regression.

Recently, Online Mirror Descent (OMD)¹ has been proposed as a general meta-algorithm for online learning, parametrized by a regularizer [16]. By appropriately choosing the regularizer, most online learning algorithms are recovered as special cases of OMD. Moreover, performance guarantees can also be derived simply by instantiating the general OMD bounds to the specific regularizer being used. So, for all the first-order online learning algorithms, it is possible to prove regret bounds of the order of $\mathcal{O}(f(\mathbf{u})\sqrt{T})$, where T is the number of rounds and $f(\mathbf{u})$ is the regularizer used in OMD, evaluated on the competitor vector \mathbf{u} . Hence, different choices of the regularizer will give rise to different algorithms and guarantees. For example, p -norm algorithms can be derived from the squared L_p -norm regularization, while EG can be derived from the entropic one. In particular for the Euclidean regularizer $\eta\sqrt{T}\|\mathbf{u}\|^2$, we have a regret bound of $\mathcal{O}(\sqrt{T}(\|\mathbf{u}\|^2/\eta + \eta))$. Knowing $\|\mathbf{u}\|$ it is possible to tune η to have a $\mathcal{O}(\|\mathbf{u}\|\sqrt{T})$ bound, that is optimal [1]. On the other hand, EG has a regret bound of $\mathcal{O}(\sqrt{T \log d})$, where d is the dimension of the space.

In this paper, I use OMD to extend EG to infinite dimensional spaces, through the use of a carefully designed time-varying regularizer. The algorithm, that I call Dimension-Free Exponentiated Gradient (DFEG), does not need direct access to single components of the vectors, rather it only requires

¹The algorithm should be more correctly called Follow the Regularized Leader, however here I follow Shalev-Shwartz in [16], and I will denote it by OMD.

to access them through inner products. Hence, DFEG can be used with kernels too, extending for the first time EG to the kernel domain. I prove a regret bound of $\mathcal{O}(\|\mathbf{u}\| \log(\|\mathbf{u}\|T + 1)\sqrt{T})$. Up to logarithmic terms, the bound of DFEG is equal to the optimal bound obtained through the knowledge of $\|\mathbf{u}\|$, but it does not require the tuning of any parameter.

I built upon ideas of [19], but I designed my new algorithm as an instantiation of OMD, rather than using an ad-hoc analysis. I believe that this route increases the comprehension of the inner working of the algorithm, its relation to other algorithms, and it makes easier to extend it in other directions as well. In order to analyze DFEG, I also introduce new and general techniques to cope with time-varying regularizers for OMD, using the *local smoothness* of the dual of the regularization function, that might be of independent interest. I also extend and improve the lower bound in [19], to match the upper bound of DFEG up to a $\sqrt{\log T}$ term, and to show an implicit trade-off on the regret versus different competitors.

1.1 Related works

Exponentiated gradient algorithms have been proposed by [9]. The algorithms have multiplicative updates and regret bounds that depend logarithmically on the dimension of the input space. In particular, they proposed a version of EG where the weights are not normalized, called EGU.

A closer algorithm to mine is the epoch-free in [19]. Indeed, DFEG is equivalent to theirs when used on one dimensional problems. However, the extension to infinite dimensional spaces is nontrivial and very different in nature from their extension to d -dimensional problems, that consists on running a copy of the algorithm independently on each coordinate. Their regret bound depends on the dimension of the space and can neither be used with infinite dimensional spaces nor with kernels.

Vovk proposed two algorithms for square loss, with regret bounds of $\mathcal{O}((\|\mathbf{u}\| + Y)\sqrt{T})$ and $\mathcal{O}(\|\mathbf{u}\|\sqrt{T})$ respectively, where Y is an upper bound on the range of the target values [20]. A matching lower bound is also presented, proving the optimality of the second algorithm. However, the algorithms seem specific to the square loss and it is not possible to adapt them to other losses. Indeed, the lower bound I prove shows that for linear and Lipschitz losses a $\sqrt{\log(\|\mathbf{u}\|T)}$ term is unavoidable. Moreover, the second algorithm, being an instantiation of the Aggregating Algorithm [21], does not seem to have an efficient implementation.

My algorithm also shares similarities in spirit with the family of self-confident algorithms [2, 7, 15], in which the algorithm self-tunes its parameters based on internal estimates.

From the point of view of the proof technique, the primal-dual analysis of OMD is due to [15, 17]. Starting from the work of [8], it is now clear that OMD can be easily analyzed using only a few basic convex duality properties. See the recent survey [16] for a lucid description of these developments. The time-varying regularization for OMD has been explored in [4, 12, 15], but in none of these works does the negative terms in the bound due to the time-varying regularizer play a decisive role. The use of the local estimates of strong smoothness is new, as far as I know. A related way to have a local analysis is through the *local norms* [16], but my approach is better tailored to my needs.

2 Problem setting and definitions

In the online learning scenario the learning algorithms work in rounds [3]. Let \mathbb{X} a Euclidean vector space², at each round t , an instance $\mathbf{x}_t \in \mathbb{X}$, is presented to the algorithm, which then predicts a label $\hat{y}_t \in \mathbb{R}$. Then, the correct label y_t is revealed, and the algorithm pays a loss $\ell(\hat{y}_t, y_t)$, for having predicted \hat{y}_t , instead of y_t . The aim of the online learning algorithm is to minimize the cumulative sum of the losses, on any sequence of data/labels $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$. Typical examples of loss functions are, for example, the absolute loss, $|\hat{y}_t - y_t|$, and the hinge loss, $\max(1 - \hat{y}_t y_t, 0)$. Note that the loss function can change over time, so in the following I will denote by $\ell_t : \mathbb{R} \rightarrow \mathbb{R}$ the generic loss function received by the algorithm at time t . In this paper I focus on linear prediction of the form $\hat{y}_t = \langle \mathbf{w}_t, \mathbf{x}_t \rangle$, where $\mathbf{w}_t \in \mathbb{X}$ represents the hypothesis of the online algorithm at time t .

²All the theorems hold also in general Hilbert spaces, but for simplicity of exposition I consider a Euclidean setting.

Algorithm 1 Dimension-Free Exponentiated Gradient.

Parameters: $0.882 \leq a \leq 1.109$, $L > 0$, $\delta > 0$.

Initialize: $\theta_1 = \mathbf{0} \in \mathbb{X}$, $H_0 = \delta$

for $t = 1, 2, \dots$ **do**

 Receive $\|\mathbf{x}_t\|$, where $\mathbf{x}_t \in \mathbb{X}$

 Set $H_t = H_{t-1} + L^2 \max(\|\mathbf{x}_t\|, \|\mathbf{x}_t\|^2)$

 Set $\alpha_t = a\sqrt{H_t}$, $\beta_t = H_t^{\frac{3}{2}}$

if $\|\theta_t\| == 0$ **then** choose $\mathbf{w}_t = \mathbf{0}$

else choose $\mathbf{w}_t = \frac{\theta_t}{\beta_t \|\theta_t\|} \exp\left(\frac{\|\theta_t\|}{\alpha_t}\right)$

 Suffer loss $\ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle)$

 Update $\theta_{t+1} = \theta_t - \partial \ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle) \mathbf{x}_t$

end for

We strive to design online learning algorithms for which it is possible to prove a relative *regret* bound. Such analysis bounds the *regret*, that is the difference between the cumulative loss of the algorithm, $\sum_{t=1}^T \ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle)$, and the one of an arbitrary and fixed competitor \mathbf{u} , $\sum_{t=1}^T \ell_t(\langle \mathbf{u}, \mathbf{x}_t \rangle)$. We will consider L -Lipschitz losses, that is $|\ell_t(y) - \ell_t(y')| \leq L|y - y'|$, $\forall y, y'$.

I now introduce some basic notions of convex analysis that are used in the paper. I refer to [14] for definitions and terminology. I consider functions $f : \mathbb{X} \rightarrow \mathbb{R}$ that are closed and convex. Given a closed and convex function f with domain $S \subseteq \mathbb{X}$, its Fenchel conjugate $f^* : \mathbb{X} \rightarrow \mathbb{R}$ is defined as $f^*(\mathbf{u}) = \sup_{\mathbf{v} \in S} (\langle \mathbf{v}, \mathbf{u} \rangle - f(\mathbf{v}))$. The Fenchel-Young inequality states that $f(\mathbf{u}) + f^*(\mathbf{v}) \geq \langle \mathbf{u}, \mathbf{v} \rangle$ for all \mathbf{v}, \mathbf{u} . A vector \mathbf{x} is a subgradient of a convex function f at \mathbf{v} if $f(\mathbf{u}) - f(\mathbf{v}) \geq \langle \mathbf{u} - \mathbf{v}, \mathbf{x} \rangle$ for any \mathbf{u} in the domain of f . The differential set of f at \mathbf{v} , denoted by $\partial f(\mathbf{v})$, is the set of all the subgradients of f at \mathbf{v} . If f is also differentiable at \mathbf{v} , then $\partial f(\mathbf{v})$ contains a single vector, denoted by $\nabla f(\mathbf{v})$, which is the gradient of f at \mathbf{v} . Strong convexity and strong smoothness are key properties in the design of online learning algorithms, they are defined as follows. A function f is γ -strongly convex with respect to a norm $\|\cdot\|$ if for any \mathbf{u}, \mathbf{v} in its domain, and any $\mathbf{x} \in \partial f(\mathbf{u})$,

$$f(\mathbf{v}) \geq f(\mathbf{u}) + \langle \mathbf{v} - \mathbf{u}, \mathbf{x} \rangle + \frac{\gamma}{2} \|\mathbf{u} - \mathbf{v}\|^2.$$

The Fenchel conjugate f^* of a γ -strongly convex function f is everywhere differentiable and $\frac{1}{\gamma}$ -strongly smooth [8], this means that for all $\mathbf{u}, \mathbf{v} \in \mathbb{X}$,

$$f^*(\mathbf{v}) \leq f^*(\mathbf{u}) + \langle \mathbf{v} - \mathbf{u}, \nabla f^*(\mathbf{u}) \rangle + \frac{1}{2\gamma} \|\mathbf{u} - \mathbf{v}\|_*^2.$$

In the remainder of the paper all the norms considered will be the L_2 ones.

3 Dimension-Free Exponentiated Gradient

In this section I describe the DFEG algorithm. The pseudo-code is in Algorithm 1. It shares some similarities with the exponentiated gradient with unnormalized weights algorithm [9], to the self-tuning variant of exponentiated gradient in [15], and to the epoch-free algorithm in [19]. However, note that it does not access to single coordinates of \mathbf{w}_t and \mathbf{x}_t , but only their inner products. Hence, we expect the algorithm not to depend on the dimension of \mathbb{X} , that can be even infinite. In other words, DFEG can be used with kernels as well, on contrary of all the mentioned algorithms above.

For the DFEG algorithm we have the following regret bound, that will be proved in Section 4.

Theorem 1. *Let $0.882 \leq a \leq 1.109$, $\delta > 0$, then, for any sequence of input vectors $\{\mathbf{x}_t\}_{t=1}^T$, any sequence of L -Lipschitz convex losses $\{\ell_t(\cdot)\}_{t=1}^T$, and any $\mathbf{u} \in \mathbb{X}$, the following bound on the regret holds for Algorithm 1*

$$\sum_{t=1}^T \ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle) - \sum_{t=1}^T \ell_t(\langle \mathbf{u}, \mathbf{x}_t \rangle) \leq \frac{4 \exp(1 + \frac{1}{a})}{L\sqrt{\delta}} + a\|\mathbf{u}\|\sqrt{H_T} \left(\ln \left(H_T^{\frac{3}{2}} \|\mathbf{u}\| \right) - 1 \right),$$

where $H_T = \delta + \sum_{t=1}^T L^2 \max(\|\mathbf{x}_t\|, \|\mathbf{x}_t\|^2)$.

The bound has a logarithmic part, typical of the family of exponentiated gradient algorithms, but instead of depending on the dimension, it depends on the norm of the competitor, $\|\mathbf{u}\|$. Hence, the regret bound of DFEG holds for infinite dimensional spaces as well, that is, it is dimension-free.

It is interesting to compare this bound with the usual bound for online learning using an L_2 regularizer. Using a time-varying regularizer $f_t(\mathbf{w}) = \frac{\sqrt{t}}{\eta} \|\mathbf{w}\|^2$ it is easy to see, e.g. [15], that the bound would be³ $\mathcal{O}((\|\mathbf{u}\|^2/\eta + \eta)\sqrt{T})$. If an upper bound U on $\|\mathbf{u}\|$ is known, we can use it to tune η to obtain an upper bound of the order of $\mathcal{O}(U\sqrt{T})$. On the other hand, we obtain for DFEG a bound of $\mathcal{O}(\|\mathbf{u}\| \log(\|\mathbf{u}\|T + 1)\sqrt{T})$, that is optimal bound, up to logarithmic terms, without knowing U . So my bound goes to constant if the norm of the competitor goes to zero. However, note that, for any fixed competitor, the gradient descent bound is asymptotically better.

The lower bound on the range of a we get comes from technical details of the analysis. The parameter a is directly linked to the leading constant of the regret bound; therefore, it is intuitive that the range of acceptable values must have a lower bound different from zero. This is also confirmed by the lower bound in Theorem 2 below.

Notice that the bound is data-dependent because it depends on the sequence of observed input vectors \mathbf{x}_t . A data-independent bound can be easily obtained from the upper bound on the norm of the input vectors. The use of the function $\max(\|\mathbf{x}_t\|, \|\mathbf{x}_t\|^2)$ is necessary to have such a data-dependent bound and it seems that it cannot be avoided in order to prove the regret bound.

It is natural to ask if the log term in the bound can be avoided. Extending Theorem 7 in [19], we can reply in the negative to this question. In particular, the following theorem shows that the regret of any online learning algorithm has to satisfy a trade-off between the guarantees against the competitor with norm equal to zero and the ones against other competitors. A similar trade-off has been proven in the expert settings [5].

Theorem 2. *Fix a non-trivial vector space \mathcal{X} , a specific online learning algorithm, and let the sequence of losses be composed by linear losses. If the algorithm guarantees a zero regret against the competitor with zero L_2 norm, then there exists a sequence of T vectors in \mathcal{X} , such that the regret against any other competitor is $\Omega(T)$. On the other hand, if the algorithm guarantees a regret at most of $\epsilon > 0$ against the competitor with zero L_2 norm, then, for any $0 < \eta < 1$, there exists a T_0 and a sequence of $T \geq T_0$ unitary norm vectors $\mathbf{z}_t \in \mathcal{X}$, and a vector $\mathbf{u} \in \mathcal{X}$ such that*

$$\sum_{t=1}^T \langle \mathbf{u}, \mathbf{z}_t \rangle - \sum_{t=1}^T \langle \mathbf{w}_t, \mathbf{z}_t \rangle \geq (1 - \eta) \|\mathbf{u}\| \sqrt{\frac{1}{\log 2}} \sqrt{T \log \left(\frac{\eta \|\mathbf{u}\| \sqrt{T}}{3\epsilon} \right)} - 2.$$

The proof can be found in the supplementary material. It is possible to show that the optimal η is of the order of $\frac{1}{\log T}$, so that the leading constant approaches $\sqrt{\frac{1}{\log 2}} \approx 1.2011$ when T goes to infinity.

It is also interesting to note that an L_2 regularizer suffers a loss of $\mathcal{O}(\sqrt{T})$ against a competitor with zero norm, that cancels the $\sqrt{\log T}$ term.

4 Analysis

In this section I prove my main result. I will first briefly introduce the general OMD algorithm with time-varying regularizers on which my algorithm is based.

4.1 Online mirror descent and local smoothness

Algorithm 2 is a generic meta-algorithm for online learning. Most of the online learning algorithms can be derived from it, choosing the functions f_t and the vectors \mathbf{z}_t . The following lemma, that is a generalization of Corollary 4 in [8], Corollary 3 in [4], and Lemma 1 in [12], is the main tool to prove the regret bound for the DFEG algorithm. The proof is in the supplementary material.

³Despite what claimed in Section 1 of [19], the use of the time-varying regularizer $f_t(\mathbf{w}) = \frac{\sqrt{t}}{\eta} \|\mathbf{w}\|^2$ guarantees a sublinear regret for unconstrained online convex optimization, for any $\eta > 0$.

Algorithm 2 Time-varying Online Mirror Descent

Parameters: A sequence of convex functions f_1, f_2, \dots defined on $S \subseteq \mathbb{X}$.
Initialize: $\theta_1 = \mathbf{0} \in \mathbb{X}$
for $t = 1, 2, \dots$ **do**
 Choose $\mathbf{w}_t = \nabla f_t^*(\theta_t)$
 Observe $\mathbf{z}_t \in \mathbb{X}$
 Update $\theta_{t+1} = \theta_t + \mathbf{z}_t$
end for

Lemma 1. Assume Algorithm 2 is run with functions f_1, f_2, \dots defined on a common domain $S \subseteq \mathbb{X}$. Then for any $\mathbf{w}'_t, \mathbf{u} \in S$ we have

$$\sum_{t=1}^T \langle \mathbf{z}_t, \mathbf{u} - \mathbf{w}'_t \rangle \leq f_T(\mathbf{u}) + \sum_{t=1}^T (f_t^*(\theta_{t+1}) - f_{t-1}^*(\theta_t) - \langle \mathbf{w}'_t, \mathbf{z}_t \rangle),$$

where we set $f_0^*(\mathbf{w}'_1) = 0$. Moreover, if f_1^*, f_2^*, \dots are twice differentiable, and $\max_{0 \leq \tau \leq 1} \|\nabla^2 f_t^*(\theta_t + \tau \mathbf{z}_t)\| \leq \lambda_t$, then we have

$$f_t^*(\theta_{t+1}) - f_{t-1}^*(\theta_t) - \langle \mathbf{w}_t, \mathbf{z}_t \rangle \leq f_t^*(\theta_t) - f_{t-1}^*(\theta_t) + \frac{\lambda_t}{2} \|\mathbf{z}_t\|^2.$$

Note that the above Lemma is usually stated assuming the strong convexity of f_t , that is equivalent to the strong smoothness of f_t^* , that in turns for twice differentiable functions is equivalent to a global bound on the norm of the Hessian of f_t^* (see Theorem 2.1.6 in [11]). Here I take a different route, assuming the functions f_t^* to be twice differentiable, but using the weaker hypothesis of local boundedness of the Hessian of f_t^* . Hence, for twice differentiable conjugate functions, this bound is always tighter than the ones in [4, 8, 12]. Indeed, in our case, the global strong smoothness cannot be used to prove any meaningful regret bound.

We derive the Dimension-Free Exponentiated Gradient from the general OMD above. Set in Algorithm 2 $f_t(\mathbf{w}) = \alpha_t \|\mathbf{w}\|(\log(\beta_t \|\mathbf{w}\|) - 1)$, where α_t and β_t are defined in Algorithm 1, and $\mathbf{z}_t = -\partial \ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle) \mathbf{x}_t$. The proof idea of my theorem is the following. First, assume that we are on a round where we have a local upper bound on the norm of the Hessian f_t^* . The usual approach in these kind of proof is to have a regularizer that is growing over time as \sqrt{t} , so that the terms $f_t^*(\theta_t) - f_{t-1}^*(\theta_t)$ are negative and can be safely discarded. At the same time the sum of the squared norms of the gradients will typically be of the order of $\mathcal{O}(\sqrt{T})$, giving us a $\mathcal{O}(\sqrt{T})$ regret bound (see for example the proofs in [4]). However, following this approach in DFEG we would have that the sum of norms of the squared gradients grows much faster than $\mathcal{O}(\sqrt{T})$. This is due to the fact that the global strong smoothness is too small. Hence I introduce a different proof method. In the following, I will show the surprising result that with my choice of the regularizers f_t , the terms $f_t^*(\theta_t) - f_{t-1}^*(\theta_t)$ and the squared norm of the gradient cancel out. Notice that already in [12, 13] it has been advocated not to discard those terms to obtain tighter bounds. Here the same terms play a major role in the proof and they are present thanks to the time-varying regularization. This is in agreement with Theorem 9 in [19] that rules out algorithms with a fixed regularizer to obtain regret bounds like Theorem 1.

It remains to bound the regret in the rounds where we do not have an upper bound on the norm of the Hessian. In these rounds I show that the norm of \mathbf{w}_t (and θ_t) is small enough so that the regret is still bounded, thanks to the choice of β_t .

4.2 Proof of the main result

We start defining the new regularizer and show its properties in the following Lemma (proof in the supplementary material). Note the similarities with EGU, where the regularizer is $\sum_{i=1}^d w_i(\log(w_i) - 1)$, $\mathbf{w} \in \mathbb{R}^d, w_i \geq 0$ [9].

Lemma 2. Define $f(\mathbf{w}) = \alpha \|\mathbf{w}\|(\ln(\beta \|\mathbf{w}\|) - 1)$, for $\alpha, \beta > 0$. The following properties hold

- $f^*(\theta) = \frac{\alpha}{\beta} \exp \frac{\|\theta\|}{\alpha}$.

- $\nabla f^*(\boldsymbol{\theta}) = \frac{\boldsymbol{\theta}}{\beta \|\boldsymbol{\theta}\|} \exp \frac{\|\boldsymbol{\theta}\|}{\alpha}$.
- $\|\nabla^2 f^*(\boldsymbol{\theta})\|_2 \leq \frac{1}{\beta \min(\|\boldsymbol{\theta}\|, \alpha)} \exp(\frac{\|\boldsymbol{\theta}\|}{\alpha})$.

Equipped with a local upper bound on the Hessian of f^* , we can now use Lemma 1. We notice that Lemma 1 also guides us in the choice of the sequences α_t . In fact if we want the regret to be $\tilde{O}(\sqrt{T})$, α_t must be $\tilde{O}(\sqrt{T})$ too.

In the proof of Theorem 1 we also use the following three technical lemmas, whose proofs are in the supplementary material. The first two are used to upper bound the exponential function with quadratic functions.

Lemma 3. *Let $M > 0$, then for any $\frac{\exp(M)}{M^2+1} \leq p \leq \exp(M)$, and $0 \leq x \leq M$, we have $\exp(x) \leq p + \frac{\exp(M)-p}{M^2}x^2$.*

Lemma 4. *Let $M > 0$, then for any $0 \leq x \leq M$, we have $\exp(x) \leq 1 + x + \frac{\exp(M)-1-M}{M^2}x^2$.*

Lemma 5. *For any $p, q > 0$ we have that $\frac{2}{\sqrt{p}} - \frac{2}{\sqrt{p+q}} \geq \frac{q}{(p+q)^{\frac{3}{2}}}$.*

Proof of Theorem 1. In the following denote by $n(\mathbf{x}) := \max(\|\mathbf{x}\|, \|\mathbf{x}\|^2)$. We will use Lemma 1 to upper bound the regret of DFEG. Hence, using the notation in Algorithm 1, set $\mathbf{z}_t = -\partial \ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle) \mathbf{x}_t$, and $f_t(\mathbf{w}) = \alpha_t \|\mathbf{w}\|(\log(\beta_t \|\mathbf{w}\|) - 1)$. Observe that, by the hypothesis on ℓ_t , we have $\|\mathbf{z}_t\| \leq L \|\mathbf{x}_t\|$. We first consider two cases, based on the norm of $\boldsymbol{\theta}_t$.

Case 1: $\|\boldsymbol{\theta}_t\| > \alpha_t + \|\mathbf{z}_t\|$.

With this assumption, and using the third property of Lemma 2, we have

$$\max_{0 \leq \tau \leq 1} \|\nabla^2 f_t^*(\boldsymbol{\theta}_t + \tau \mathbf{z}_t)\| \leq \max_{0 \leq \tau \leq 1} \frac{\exp\left(\frac{\|\boldsymbol{\theta}_t + \tau \mathbf{z}_t\|}{\alpha_t}\right)}{\beta_t \min(\|\boldsymbol{\theta}_t + \tau \mathbf{z}_t\|, \alpha_t)} \leq \frac{\exp\left(\frac{\|\boldsymbol{\theta}_t\| + \|\mathbf{z}_t\|}{\alpha_t}\right)}{\beta_t \alpha_t}.$$

We now use the second statement of Lemma 1. We have that $\frac{\lambda_t \|\mathbf{z}_t\|^2}{2} + f_t^*(\boldsymbol{\theta}_t) - f_{t-1}^*(\boldsymbol{\theta}_t)$ can be upper bounded by

$$\begin{aligned} & \frac{\|\mathbf{z}_t\|^2}{2\alpha_t\beta_t} \exp\left(\frac{\|\boldsymbol{\theta}_t\| + \|\mathbf{z}_t\|}{\alpha_t}\right) + \frac{\alpha_t}{\beta_t} \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) - \frac{\alpha_{t-1}}{\beta_{t-1}} \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_{t-1}}\right) \\ & \leq \frac{\|\mathbf{z}_t\|^2}{2\alpha_t\beta_t} \exp\left(\frac{\|\boldsymbol{\theta}_t\| + \|\mathbf{z}_t\|}{\alpha_t}\right) + \frac{\alpha_t}{\beta_t} \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) - \frac{\alpha_{t-1}}{\beta_{t-1}} \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \\ & = \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \left(\frac{\|\mathbf{z}_t\|^2}{2aH_t^2} \exp\left(\frac{\|\mathbf{z}_t\|}{\alpha_t}\right) + \frac{a}{H_t} - \frac{a}{H_{t-1}} \right). \end{aligned} \quad (1)$$

We will now prove that the term in the parenthesis of (1) is negative. It can be rewritten as

$$\frac{\|\mathbf{z}_t\|^2}{2aH_t^2} \exp\left(\frac{\|\mathbf{z}_t\|}{\alpha_t}\right) + \frac{a}{H_t} - \frac{a}{H_{t-1}} = \frac{\|\mathbf{z}_t\|^2 H_{t-1} \exp\left(\frac{\|\mathbf{z}_t\|}{\alpha_t}\right) - 2a^2 H_{t-1} L^2 n(\mathbf{x}_t) - 2a^2 L^4 (n(\mathbf{x}_t))^2}{2aH_t^2 H_{t-1}},$$

and from the expression of α_t we have that $\frac{\|\mathbf{z}_t\|}{\alpha_t} \leq \frac{1}{a}$, so we now use Lemma 3 with $p = 2a^2$ and $M = 1/a$. These are valid settings because $\frac{\exp(\frac{1}{a})}{\frac{1}{a^2}+1} \leq 2a^2 \leq \exp(\frac{1}{a})$, $\forall 0.825 \leq a \leq 1.109$, as it can be verified numerically.

$$\begin{aligned} & \frac{\|\mathbf{z}_t\|^2}{2aH_t^2} \exp\left(\frac{\|\mathbf{z}_t\|}{\alpha_t}\right) + \frac{a}{H_t} - \frac{a}{H_{t-1}} \\ & \leq \frac{\|\mathbf{z}_t\|^2 H_{t-1} \left(2a^2 + a^2(\exp(\frac{1}{a}) - 2a^2) \frac{\|\mathbf{z}_t\|^2}{\alpha_t^2}\right) - 2a^2 H_{t-1} L^2 n(\mathbf{x}_t) - 2a^2 L^4 (n(\mathbf{x}_t))^2}{2aH_t^2 H_{t-1}} \\ & \leq \frac{L^2 \|\mathbf{x}_t\|^2 H_{t-1} \left(2a^2 + a^2(\exp(\frac{1}{a}) - 2a^2) \frac{L^2 \|\mathbf{x}_t\|^2}{a^2 H_t}\right) - 2a^2 H_{t-1} L^2 \|\mathbf{x}_t\|^2 - 2a^2 L^4 \|\mathbf{x}_t\|^2}{2aH_t^2 H_{t-1}} \\ & \leq \frac{L^4 \|\mathbf{x}_t\|^4 (\exp(\frac{1}{a}) - 4a^2)}{2aH_t^2 H_{t-1}} \leq 0, \end{aligned} \quad (2)$$

where in last step we used the fact that $\exp(\frac{1}{a}) \leq 4a^2, \forall a \geq 0.882$, as again it can be verified numerically.

Case 2: $\|\boldsymbol{\theta}_t\| \leq \alpha_t + \|\mathbf{z}_t\|$.

We use the first statement of Lemma 1, setting $\mathbf{w}'_t = \mathbf{w}_t$ if $\|\boldsymbol{\theta}\| \neq 0$, and $\mathbf{w}'_t = \mathbf{0}$ otherwise. In this way, from the second property of Lemma 2, we have that $\|\mathbf{w}'_t\| \leq \frac{1}{\beta_t} \exp(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t})$. Note that any other choice of \mathbf{w}'_t satisfying the the previous relation on the norm of \mathbf{w}'_t would have worked as well.

$$\begin{aligned} f_t^*(\boldsymbol{\theta}_{t+1}) - f_{t-1}^*(\boldsymbol{\theta}_t) &= \frac{\alpha_t}{\beta_t} \exp\left(\frac{\|\boldsymbol{\theta}_{t+1}\|}{\alpha_t}\right) - \frac{\alpha_{t-1}}{\beta_{t-1}} \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_{t-1}}\right) \\ &\leq \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \left(\frac{\alpha_t}{\beta_t} \exp\left(\frac{\|\mathbf{z}_t\|}{\alpha_t}\right) - \frac{\alpha_{t-1}}{\beta_{t-1}}\right) = a \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{\exp\left(\frac{\|\mathbf{z}_t\|}{a\sqrt{H_t}}\right) H_{t-1} - H_t}{H_{t-1} H_t}. \end{aligned} \quad (3)$$

Remembering that $\frac{\|\mathbf{z}_t\|}{\alpha_t} \leq \frac{1}{a}$, and using Lemma 4 with $M = \frac{1}{a}$, we have

$$\begin{aligned} H_{t-1} \exp\left(\frac{\|\mathbf{z}_t\|}{a\sqrt{H_t}}\right) - H_{t-1} - L^2 n(\mathbf{x}_t) &\leq H_{t-1} \exp\frac{L\|\mathbf{x}_t\|}{a\sqrt{H_t}} - H_{t-1} - L^2 \|\mathbf{x}_t\|^2 \\ &\leq H_{t-1} \left(1 + \frac{L\|\mathbf{x}_t\|}{a\sqrt{H_t}} + a^2 \left(\exp\left(\frac{1}{a}\right) - 1 - \frac{1}{a}\right) \frac{L^2 \|\mathbf{x}_t\|^2}{a^2 H_t}\right) - H_{t-1} - L^2 \|\mathbf{x}_t\|^2 \\ &= \frac{LH_{t-1}\|\mathbf{x}_t\|}{a\sqrt{H_t}} + \left(\exp\left(\frac{1}{a}\right) - 1 - \frac{1}{a}\right) \frac{L^2 H_{t-1} \|\mathbf{x}_t\|^2}{H_t} - L^2 \|\mathbf{x}_t\|^2 \\ &\leq \frac{LH_{t-1}\|\mathbf{x}_t\|}{a\sqrt{H_t}} + L^2 \|\mathbf{x}_t\|^2 \left(\exp\left(\frac{1}{a}\right) - 2 - \frac{1}{a}\right) \leq \frac{LH_{t-1}\|\mathbf{x}_t\|}{a\sqrt{H_t}}, \end{aligned} \quad (4)$$

where in the last step we used the fact that $\exp(\frac{1}{a}) - 2 - \frac{1}{a} \leq 0, \forall a \geq 0.873$, verified numerically. Putting together (3) and (4), we have

$$\begin{aligned} f_t^*(\boldsymbol{\theta}_{t+1}) - f_{t-1}^*(\boldsymbol{\theta}_t) - \langle \mathbf{w}'_t, \mathbf{z}_t \rangle &\leq \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{L\|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}} - \langle \mathbf{w}'_t, \mathbf{z}_t \rangle \\ &\leq \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{L\|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}} + L\|\mathbf{w}'_t\| \|\mathbf{x}_t\| \leq \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{L\|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}} + \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{L\|\mathbf{x}_t\|}{\beta_t} \\ &= 2 \exp\left(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}\right) \frac{L\|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}} \leq \frac{2 \exp(1 + \frac{1}{a}) L \|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}}, \end{aligned} \quad (5)$$

where in the second inequality we used the Cauchy-Schwarz inequality and the Lipschitzness of ℓ_t , in the third the bound on the norm of \mathbf{w}'_t , and in the last inequality the fact that $\|\boldsymbol{\theta}_t\| \leq \alpha_t + \|\mathbf{z}_t\|$ implies $\exp(\frac{\|\boldsymbol{\theta}_t\|}{\alpha_t}) \leq \exp(1 + \frac{1}{a})$. Putting together (2) and (5) and summing over t , we have

$$\begin{aligned} \sum_{t=1}^T (f_t^*(\boldsymbol{\theta}_{t+1}) - f_{t-1}^*(\boldsymbol{\theta}_t) - \langle \mathbf{w}'_t, \mathbf{z}_t \rangle) &\leq \sum_{t=1}^T \frac{2 \exp(1 + \frac{1}{a}) L \|\mathbf{x}_t\|}{H_t^{\frac{3}{2}}} \leq \frac{2}{L} \sum_{t=1}^T \frac{\exp(1 + \frac{1}{a}) L^2 \|\mathbf{x}_t\|}{(\sum_{j=1}^t L^2 \|\mathbf{x}_j\| + \delta)^{\frac{3}{2}}} \\ &\leq \frac{4 \exp(1 + \frac{1}{a})}{L} \sum_{t=1}^T \left(\frac{1}{\sqrt{\sum_{j=1}^{t-1} L^2 \|\mathbf{x}_j\| + \delta}} - \frac{1}{\sqrt{\sum_{j=1}^t L^2 \|\mathbf{x}_j\| + \delta}} \right) \leq \frac{4 \exp(1 + \frac{1}{a})}{L\sqrt{\delta}}, \end{aligned}$$

where in the third inequality we used Lemma 5.

The stated bound can be obtained observing that $\ell_t(\langle \mathbf{w}_t, \mathbf{x}_t \rangle) - \ell_t(\langle \mathbf{u}, \mathbf{x}_t \rangle) \leq \langle \mathbf{u} - \mathbf{w}_t, \mathbf{z}_t \rangle$, from the convexity of ℓ_t and the definition of \mathbf{z}_t . \square

5 Experiments

A full empirical evaluation of DFEG is beyond the scope of this paper. Here I just want to show the empirical effect of some of its theoretical properties. In all the experiments I used the absolute loss,




Figure 1: Left: regret versus number of input vectors on synthetic dataset. Center and Right: total loss for DFEG and Kernel GD on the *cadata* and *cpusmall* dataset respectively.

so $L = 1$, a is set to the minimal value allowed by Theorem 1 and $\delta = 1$. I denote by Kernel GD the OMD with the regularizer $\frac{\sqrt{t}}{\eta} \|w\|^2$.

First, I generated synthetic data as in the proof of Theorem 2, that is the input vectors are all the same and the y_t is equal to 1 for the t even and -1 for the others. In this case we know that the optimal predictor has norm equal to zero and we can exactly calculate the value of the regret. Figure 1(left) I have plotted the regret as a function of the number of input vectors. As predicted by the theory, DFEG has a constant regret, while Kernel GD has a regret of the form $\mathcal{O}(\eta\sqrt{T})$. Hence, it can have a constant regret only when η is set to zero, and this can be done only with prior knowledge of $\|u\|$, that is impossible in practical applications.

For the second experiment, I analyzed the behavior of DFEG on two real word regression datasets, *cadata* and *cpusmall*⁴. I used the Gaussian kernel with variance equal to the average distance between training input vectors. I have plotted in Figure 1(center) the final cumulative loss of DFEG and the ones of GD with varying values of η . We see that, while the performance of Kernel GD can be better of the one of DFEG, as predicted by the theory, its performance varies greatly in relation to η . On the other hand the performance of DFEG is close to the optimal one without the need to tune any parameters. It is also worth noting the catastrophic result we can get from a wrong tuning of η in GD. Similar considerations hold for the *cpusmall* dataset in Figure 1(right).

6 Discussion

I have presented a new algorithm for online learning, the first one in the family of exponentiated gradient to be dimension-free. Thanks to new analysis tools, I have proved that DFEG attains a regret bound of $\mathcal{O}(U \log(UT + 1))\sqrt{T}$, without any parameter to tune. I also proved a lower bound that shows that the algorithm is optimal up to $\sqrt{\log T}$ term for linear and Lipschitz losses.

The problem of deriving a regret bound that depends on the sequence of the gradients, rather than on the x_t , remains open. Resolving this issue would result in the tighter $\mathcal{O}(\sqrt{\sum_{t=1}^T \ell_t(\langle w_t, x_t \rangle)})$ regret bounds in the case that the ℓ_t are smooth [18]. The difficulty in proving these kind of bounds seem to lie in the fact that (2) is negative only because $H_t - H_{t-1}$ is bigger than $\|z_t\|^2$.

Acknowledgments

I am thankful to Jennifer Batt for her help and support during the writing of this paper, to Nicolò Cesa-Bianchi for the useful comments on an early version of this work, and to Tamir Hazan for his writing style suggestions. I also thank the anonymous reviewers for their precise comments, which helped me to improve the clarity of this manuscript.

⁴<http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/>

References

- [1] J. Abernethy, A. Agarwal, P. L. Bartlett, and A. Rakhlin. A stochastic view of optimal regret through minimax duality. In *COLT*, 2009.
- [2] P. Auer, N. Cesa-Bianchi, and C. Gentile. Adaptive and self-confident on-line learning algorithms. *J. Comput. Syst. Sci.*, 64(1):48–75, 2002.
- [3] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [4] J. C. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- [5] E. Even-Dar, M. Kearns, Y. Mansour, and J. Wortman. Regret to the best vs. regret to the average. In N. H. Bshouty and C. Gentile, editors, *COLT*, volume 4539 of *Lecture Notes in Computer Science*, pages 233–247. Springer, 2007.
- [6] Y. Freund and R. E. Schapire. Large margin classification using the Perceptron algorithm. *Machine Learning*, pages 277–296, 1999.
- [7] C. Gentile. The robustness of the p-norm algorithms. *Machine Learning*, 53(3):265–299, 2003.
- [8] S. M. Kakade, S. Shalev-Shwartz, and A. Tewari. Regularization techniques for learning with matrices. *CoRR*, abs/0910.0610, 2009.
- [9] J. Kivinen and M. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132(1):1–63, January 1997.
- [10] N. Littlestone. Learning quickly when irrelevant attributes abound: a new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1988.
- [11] Y. Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer, 2003.
- [12] F. Orabona and K. Crammer. New adaptive algorithms for online classification. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 1840–1848. 2010.
- [13] F. Orabona, K. Crammer, and N. Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression, 2013. arXiv:1304.2994.
- [14] R. T. Rockafellar. *Convex Analysis (Princeton Mathematical Series)*. Princeton University Press, 1970.
- [15] S. Shalev-Shwartz. Online learning: Theory, algorithms, and applications. Technical report, The Hebrew University, 2007. PhD thesis.
- [16] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2), 2012.
- [17] S. Shalev-Shwartz and Y. Singer. A primal-dual perspective of online learning algorithms. *Machine Learning Journal*, 2007.
- [18] N. Srebro, K. Sridharan, and A. Tewari. Smoothness, low noise and fast rates. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2199–2207. 2010.
- [19] M. Streeter and B. McMahan. No-regret algorithms for unconstrained online convex optimization. In P. Bartlett, F.C.N. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2411–2419. 2012.
- [20] V. Vovk. On-line regression competitive with reproducing kernel hilbert spaces. In Jin-Yi Cai, S. Barry Cooper, and Angsheng Li, editors, *Theory and Applications of Models of Computation*, volume 3959 of *Lecture Notes in Computer Science*, pages 452–463. Springer Berlin Heidelberg, 2006.
- [21] V. G. Vovk. Aggregating strategies. In *COLT*, pages 371–386, 1990.