

Scale-Invariant Unconstrained Online Learning

Wojciech Kotłowski

Poznań University of Technology, Poland

WKOTLOWSKI@CS.PUT.POZNAN.PL

Editors: Steve Hanneke and Lev Reyzin

Abstract

We consider a variant of online convex optimization in which both the instances (input vectors) and the comparator (weight vector) are unconstrained. We exploit a natural scale invariance symmetry in our unconstrained setting: the predictions of the optimal comparator are invariant under any linear transformation of the instances. Our goal is to design online algorithms which also enjoy this property, i.e. are scale-invariant. We start with the case of coordinate-wise invariance, in which the individual coordinates (features) can be arbitrarily rescaled. We give an algorithm, which achieves essentially optimal regret bound in this setup, expressed by means of a coordinate-wise scale-invariant norm of the comparator. We then study general invariance with respect to arbitrary linear transformations. We first give a negative result, showing that no algorithm can achieve a meaningful bound in terms of scale-invariant norm of the comparator in the worst case. Next, we complement this result with a positive one, providing an algorithm which “almost” achieves the desired bound, incurring only a logarithmic overhead in terms of the norm of the instances.

Keywords: Online learning, online convex optimization, scale invariance, unconstrained online learning, linear classification, regret bound.

1. Introduction

We consider the following variant of online convex optimization (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2011; Hazan, 2015). In trials $t = 1, \dots, T$, the algorithm receives an instance $\mathbf{x}_t \in \mathbb{R}^d$, on which it predicts $\hat{y}_t = \mathbf{x}_t^\top \mathbf{w}_t$ by means of a weight vector $\mathbf{w}_t \in \mathbb{R}^d$. Then, the true label y_t is revealed and the algorithm suffers loss $\ell(y_t, \hat{y}_t)$, convex in \hat{y}_t . The algorithm’s performance is evaluated by means of *regret*, the difference between the algorithm’s cumulative loss and the cumulative loss of a prediction sequence produced by a fixed comparator (weight vector) $\mathbf{u} \in \mathbb{R}^d$. The goal of the algorithm is to minimize its regret for every data sequence $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$ and every comparator \mathbf{u} . This framework includes numerous machine learning scenarios, such as linear classification (with convex surrogate losses) and regression.

Most of the work in online convex optimization assumes that the instances and the comparator are constrained to some bounded convex sets, often known to the algorithm in advance. In practice, however, such boundedness assumptions are often unjustified: the learner has little prior knowledge on the potential magnitude of instances, while the prior knowledge on the upper bound of the comparator seems even less realistic. Therefore, there has been much work recently dedicated to relaxing some of these prior assumptions (Streeter and McMahan, 2012; Orabona, 2013; McMahan and Abernethy, 2013; McMahan and Orabona, 2014; Orabona and Pál, 2015, 2016; Luo et al., 2016). Here, we go a step

further, dropping these assumptions entirely and treating the instances, the comparator, as well as comparator’s predictions as unconstrained.

In this paper, we exploit a natural *scale invariance* symmetry of the unconstrained setting: if we transform all instances by any invertible linear transformation \mathbf{A} , $\mathbf{x} \mapsto \mathbf{Ax}$, and simultaneously transform the comparator by the (transposed) inverse of \mathbf{A} , $\mathbf{u} \mapsto \mathbf{A}^{-\top} \mathbf{u}$, the predictions, and hence the comparator’s loss will not change. This means that the predictions of the *optimal* (loss-minimizing) comparator (if exists) are invariant under any linear transformation of the instances, so that the scale of the weight vector is only relative to the scale of the instances. Our goal is to design online algorithms which also enjoy this property, i.e. their predictions are invariant under any rescaling of the instances.

Since in the absence of any constraints, the adversary can inflict arbitrarily large regret in just one trial by choosing the instance an/or comparator sufficiently large, the regret can only be bounded by a data-dependent function $\Psi(\mathbf{u}, \{\mathbf{x}_t\}_{t=1}^T)$, which can be thought of as a penalty for the adversary for having played with sequence $\{\mathbf{x}_t\}_{t=1}^T$ and the comparator \mathbf{u} . We incorporate the scale invariance into this framework by working with Ψ which depends on the data and the comparator only through the *predictions* of \mathbf{u} . As we will see, designing the online algorithms to have their regret bounded by such Ψ will automatically lead to scale-invariant methods.

We first consider a specific form of scale-invariance, which we call *coordinate-wise invariance*, in which the individual instance coordinates (“features”) can be arbitrarily rescaled (which corresponds to choosing transformation \mathbf{A} which is diagonal). One can think of such rescaling as the change of units in which coordinates are expressed. Inspired by the work of [Ross et al. \(2013\)](#), we choose the penalty function to capture the coordinate-wise invariance in the following decomposable form:

$$\Psi(\mathbf{u}, \{\mathbf{x}_t\}_{t=1}^T) = \sum_{i=1}^d f(|u_i|_{s_{T,i}}),$$

where $s_{T,i} = \sqrt{\sum_{t=1}^T x_{t,i}^2}$ are “standard deviations”¹ of individual coordinates (so that $|u_i|_{s_{T,i}}$ measures the scale of i -th coordinate relative to comparator’s weight) and $f(x) = x\sqrt{\log(1+x^2)}$. This particular choice of f is motivated by a lower bound of [Streeter and McMahan \(2012\)](#), which indicates that such dependency is the best we can hope for. The main result of Section 3 is a scale-invariant algorithm which achieves this bound up to $O(\log T)$ factor. The algorithm is a first-order method and runs in $O(d)$ time per trial. We note that when the Euclidean norms of instances and comparator are bounded by X and U , respectively, our bound reduces to Online Gradient Descent bound of $O(UX\sqrt{T})$ ([Zinkevich, 2003](#)) up to a logarithmic factor.

We then turn to a general setup in which the instances can be rescaled by arbitrary linear transformations. A natural and analytically tractable choice is to parameterize the bound by means of a sum of squared predictions:

$$\Psi(\mathbf{u}, \{\mathbf{x}_t\}_{t=1}^T) = f(\|\mathbf{u}\|_{\mathbf{S}_T}), \quad \text{where } \|\mathbf{u}\|_{\mathbf{S}_T} \stackrel{\text{def}}{=} \sqrt{\mathbf{u}^\top \mathbf{S}_T \mathbf{u}} = \sqrt{\sum_{t=1}^T (\mathbf{x}_t^\top \mathbf{u})^2},$$

1. Throughout the paper, the terms “standard deviation”, “variance” and “covariance matrix” are used informally (the coordinate values are not shifted by their mean), to capture the scale of the instances.

and where $\mathbf{S}_T = \sum_t \mathbf{x}_t \mathbf{x}_t^\top$ is the empirical ‘‘covariance’’ matrix, and $f(x) = x\sqrt{\log(1+x^2)}$ as before. Our first result is a negative one: any algorithm can be forced by an adversary to have a regret at least $\Omega(\|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T})$ already for $d = 2$ dimensional inputs. It turns out that such a bound is meaningless, as a trivial algorithm which always predicts zero has its regret bounded by $O(\|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T})$.

Is is then the end of the story? While the above result suggests that the adversary has too much power and every algorithm fails in this setting, we show that this view is too pessimistic, complementing the negative result with a positive one. In Section 4 we derive a scale-invariant algorithm that is capable of *almost* achieving the bound expressed by Ψ above, with only a logarithmic dependence on the norm of the instances. The algorithm is a second-order method and runs in $O(d^2)$ time per trial.

1.1. Related work

A standard setup in online convex optimization (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2011; Hazan, 2015) assumes that both the instances² and the comparator are constrained to some bounded convex sets, known to the learner in advance. A recent series of papers has explored a setting in which the comparator is unconstrained and the learner needs to adapt to an unknown comparator norm (Streeter and McMahan, 2012; Orabona, 2013; McMahan and Abernethy, 2013; McMahan and Orabona, 2014; Orabona, 2014; Orabona and Pál, 2016). Most of these papers (exception being Orabona (2013)), however, assume that the loss gradients (and thus instances in our setup) are bounded. Moreover, none of these papers concerns scale-invariance.

Scale-invariant online algorithms were studied by Ross et al. (2013), who consider a setup similar to our coordinate-wise case. They, however, make a strong assumption that all individual feature constituents of the comparator predictions are bounded: $|u_i x_{t,i}| \leq C$ for all $i = 1, \dots, d$ and $t = 1, \dots, T$, where C is known to the learner. Their algorithm has a bound which depends on $\frac{C}{b_{T,i}}$, $i = 1, \dots, d$ (where $b_{t,i} = \max_{q=1, \dots, t} |x_{q,i}|$), which is in fact the worst-case upper bound on u_i ; furthermore their bound also depend as on the ratios $\frac{b_{T,i}}{x_{t_i,i}}$ (t_i being the first trial in which the i -th feature $x_{t_i,i}$ is non-zero), which can be made arbitrarily large in the worst case. Orabona et al. (2015) study a similar setup, giving a bound in terms of the quantities $C \frac{|x_{t,i}|}{b_{t,i}}$, and the algorithm still requires to know C to tune its learning rate. In contrast, we do not make any assumptions on the predictions of \mathbf{u} , and our bound depends on the actual values of u_i , solely by means of $\sqrt{u_i^2 \sum_t x_{t,i}^2}$, $i = 1, \dots, d$.

Luo et al. (2016) consider a setup similar to our full scale-invariant case, but they require an additional constraint that $|\mathbf{u}^\top \mathbf{x}_t| \leq C$ for all t , which we avoid in this work. Finally, Orabona and Pál (2015) consider a different notion of invariance, unrelated to this setup.

2. Problem Setup

We consider a variant of online convex optimization summarized in Figure 1. In each trial $t = 1, \dots, T$, an instance $\mathbf{x}_t \in \mathbb{R}^d$ is presented to the learner, which produces a weight

2. Most papers assume the bound on the (sub)gradient of loss with respect to \mathbf{w} , which translates to bound on the instances, as $\nabla_{\mathbf{w}} \ell(y, \mathbf{x}^\top \mathbf{w}) = \partial_{\hat{y}} \ell(y, \hat{y}) \cdot \mathbf{x}$.

At trial $t = 1 \dots T$:
 Instance $\mathbf{x}_t \in \mathbb{R}^d$ is revealed to the learner.
 Learner predicts with $\hat{y}_t = \mathbf{x}_t^\top \mathbf{w}_t$ for some $\mathbf{w}_t \in \mathbb{R}^d$.
 Adversary reveals label $y_t \in \mathbb{R}$.
 Learner suffers loss $\ell(y_t, \hat{y}_t)$.

Figure 1: Online learning protocol considered in this work.

vector $\mathbf{w}_t \in \mathbb{R}^d$ (possibly depending on \mathbf{x}_t) and prediction $\hat{y}_t = \mathbf{x}_t^\top \mathbf{w}_t$. Then, the true label y_t is revealed, and the learner suffers loss $\ell(y_t, \hat{y}_t)$. We assume the loss is convex in its second argument and L -Lipschitz (where L is known to the learner), i.e. subderivatives of ℓ are bounded, $|\partial_{\hat{y}} \ell(y, \hat{y})| \leq L$ for all y, \hat{y} . Two popular loss functions which fall into this framework (with $L = 1$) are logistic loss $\ell(y, \hat{y}) = \log(1 + \exp(-y\hat{y}))$, and hinge loss $\ell(y, \hat{y}) = (1 - y\hat{y})_+$. Throughout the rest of the paper, we assume $L = 1$ without loss of generality.

The performance of the learner is evaluated by means of *regret*:

$$\text{regret}_T(\mathbf{u}) = \sum_{t=1}^T \ell(y_t, \mathbf{x}_t^\top \mathbf{w}_t) - \sum_{t=1}^T \ell(y_t, \mathbf{x}_t^\top \mathbf{u}),$$

where $\mathbf{u} \in \mathbb{R}^d$ is a fixed comparator weight vector, and the dependence on the data sequence has been omitted on the left-hand side as clear from the context. The goal of the learner is to minimize its regret for every data sequence $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$ and every comparator vector \mathbf{u} .

We use the “gradient trick” (Kivinen and Warmuth, 1997; Shalev-Shwartz, 2011), which exploits the convexity of ℓ to bound $\ell(y_t, \hat{y}_t') \geq \ell(y_t, \hat{y}_t) + \partial_{\hat{y}_t} \ell(y_t, \hat{y}_t)(\hat{y}_t' - \hat{y}_t)$ for any subderivative $\partial_{\hat{y}_t} \ell(y_t, \hat{y}_t)$ at \hat{y}_t . Using this inequality in each trial with $\hat{y}_t' = \mathbf{u}^\top \mathbf{w}_t$, we get:

$$\text{regret}_T(\mathbf{u}) \leq \sum_{t=1}^T g_t \mathbf{x}_t^\top (\mathbf{w}_t - \mathbf{u}), \quad (1)$$

where we denoted the subderivative by g_t . Throughout the rest of the paper, we will only be concerned with bounding the right-hand side of (1), i.e. we will treat the loss to be linear in the prediction, $g_t \hat{y}$, with $|g_t| \leq 1$ (which follows from 1-Lipschitzness of the loss).

In this paper, contrary to previous work, we do not impose any constraints on the instances \mathbf{x}_t or the comparator \mathbf{u} , neither on the predictions $\mathbf{x}_t^\top \mathbf{u}$. Since in the absence of any constraints, the adversary can inflict arbitrarily large regret in just one trial, the regret can only be bounded by a data dependent function $\Psi(\mathbf{u}, \{\mathbf{x}_t\}_{t=1}^T)$, which we henceforth concisely denote by $\Psi_T(\mathbf{u})$, dropping the dependence on data as clear from the context. An alternative view, which will turn out to be useful, is to study *penalized regret* $\text{regret}_T(\mathbf{u}) - \Psi_T(\mathbf{u})$, i.e. the regret offset by $\Psi_T(\mathbf{u})$, where the latter can now be treated as the penalty for the adversary (a related quantity was called *benchmark* by McMahan and Abernethy (2013)). We will design online learning algorithms which aim at minimizing the penalized regret, and this will immediately imply a data-dependent regret bound expressed by $\Psi_T(\mathbf{u})$.

As in the unconstrained setup, the predictions of the optimal comparator are invariant under any linear transformation of the instances, our goal will be to design online learning

algorithms which also enjoy this property, i.e. their predictions do not change under linear transformation of the instances. As we will see, the invariance of learning algorithms will follow from an appropriate choice of the penalty function. In Section 3, we consider algorithms invariant with respect to coordinate-wise transformations. We then move to full scale invariance (for arbitrary linear transformations) in Section 4.

3. Coordinate-wise Scale Invariance

In this section we consider algorithms which are invariant under any rescaling of individual features: if we apply any coordinate-wise transformation $x_{t,i} \mapsto a_i x_{t,i}$ for some $a_i > 0$, $i = 1, \dots, d$, $t = 1, \dots, T$, the predictions of the algorithm should remain the same. Such transformation has a natural interpretation as a change of units in which the instances are measured on each coordinate. The key element is the right choice of the penalty function $\Psi_T(\mathbf{u})$, which translates into the desired bound on the regret: the penalty function should be invariant under any feature scaling, offset by the corresponding rescaling of the comparator. Inspired by Ross et al. (2013), we consider the following function which has such a property:

$$\Psi_T(\mathbf{u}) = \sum_{i=1}^d f(|u_i|_{s_{T,i}}), \quad \text{where } s_{T,i} = \sqrt{\sum_{t=1}^T x_{t,i}^2}.$$

Quantities $\{s_{T,i}\}_{i=1}^d$ (“standard deviations”) have a very natural interpretation as measuring the scale of individual features, so that $|u_i|_{s_{T,i}}$ measures the scale relative to the comparator’s weight. It remains to choose the right form of f . To this end, we use a lower bound obtained by Streeter and McMahan (2012) (translated to our setup):

Lemma 1 (Streeter and McMahan, 2012, Theorem 7) *Consider the online protocol in Figure 1 with $d = 1$ and $x_t = 1$ for all t . Choose any learning algorithm which guarantees constant regret against comparator $u = 0$. Then for any comparator $u \in \mathbb{R}$, there exists a gradient sequences $\{g_t\}_{t=1}^T$ for which $\sum_t g_t x_t (w_t - u) = \Omega\left(|u| \sqrt{T} \sqrt{\log(|u| \sqrt{T})}\right)$.*

Since $x_t = 1$ for all t , we have $s_{T,1} = \sqrt{T}$, and the theorem suggests that the best dependence on $|u_i|_{s_{T,i}}$ one can hope for is $f(x) = x \sqrt{\log x}$. This motivates us to study the function of the form (McMahan and Orabona, 2014; Orabona and Pál, 2016):

$$f(x) = x \sqrt{\alpha \log(1 + \alpha \beta^2 x^2)}, \quad (2)$$

for some $\alpha, \beta > 0$. This particular choice of parameterization will simplify the forthcoming analysis. Ross et al. (2013) have shown that if the learner knew the comparator and standard deviations of each feature in hindsight, the optimal tuning of learning rate would result in a regret bound $\sum_i |u_i|_{s_{T,i}}$ (for $g_t \in \{-1, 1\}$). We will show that without any such prior knowledge, we will be able to essentially (up to $\log(T)$ factor) achieve a bound of $\sum_i f(|u_i|_{s_{T,i}})$, incurring only a logarithmic overhead for not knowing the scale of the instances and the comparator.

Note that the problem decomposes coordinate-wise into d one-dimensional problems, as:

$$\text{regret}_T(\mathbf{u}) - \Psi_T(\mathbf{u}) = \sum_{i=1}^d \underbrace{\left(\sum_t g_t x_{t,i} (w_{t,i} - u_i) - f(|u_i|_{s_{T,i}}) \right)}_{\text{penalized regret in 1-dim problem}}.$$

Thus, it suffices to separately analyze each such one-dimensional problem, and the final bound will be obtained by summing the individual bounds for each coordinate.

3.1. Motivation

Fix $i \in \{1, \dots, d\}$ and let us temporarily drop index i for the sake of clarity. Our goal is to design an algorithm which minimizes the one dimensional penalized regret:

$$\sum_t g_t x_t (w_t - u) - f(|u|_{s_T}),$$

where $s_T = \sqrt{\sum_t x_t^2}$ and f is given by (2). If we denote $h_t = -\sum_{q \leq t} x_q g_q$, we can rewrite the penalized regret by:

$$\sum_t g_t x_t w_t + (u h_T - f(|u|_{s_T})) \leq \sum_t g_t x_t w_t + \sup_{u \geq 0} (u |h_T| - f(|u|_{s_T})),$$

where we observed that the worst-case u will have the same sign as h_T . We now use some simple facts on Fenchel duality (Boyd and Vandenberghe, 2004). Given a function $f: X \rightarrow \mathbb{R}$, $X \subseteq \mathbb{R}$, its *Fenchel conjugate* is $f^*(\theta) = \sup_{x \in X} \{\theta x - f(x)\}$. If $g(x) = f(ax)$ for some $a > 0$, then $g^*(\theta) = f^*(\theta/a)$. Choosing $X = [0, \infty)$, and $a = s_T$, we get that:³

$$\sum_t g_t x_t (w_t - u) - f(|u|_{s_T}) \leq \sum_t g_t x_t w_t + f^*\left(\frac{|h_T|}{s_T}\right). \quad (3)$$

We now use Lemma 18 by Orabona and Pál (2016) (modified to our needs):

Lemma 2 (Orabona and Pál, 2016) *Let $f(x) = x\sqrt{\alpha \log(1 + \alpha\beta^2 x^2)}$ for $\alpha, \beta > 0$ and $x \geq 0$. Then $f^*(\theta) \leq \frac{1}{\beta} e^{\frac{\theta^2}{2\alpha}}$.*

Applying Lemma 2 to (3) results in:

$$\sum_t g_t x_t (w_t - u) - f(|u|_{s_T}) \leq \sum_t g_t x_t w_t + \frac{1}{\beta} \exp\left(\frac{h_T^2}{2\alpha s_T^2}\right). \quad (4)$$

The main advantage of this bound is the elimination of unknown comparator u . We can now design learning algorithm to directly minimize the right-hand side of (4) over the worst-case choice of the data. What we derived here is essentially a variant of “regret-reward duality” (Streeter and McMahan, 2012; McMahan and Orabona, 2014; Orabona and Pál, 2016).

3. In the excluded case $s_T = 0$ the regret is trivially zero as $x_t = 0$ for all t .

Algorithm 1: Coordinate-wise scale invariant algorithm

Parameter : $\alpha > \frac{9}{8} = 1.125$.
Initialization: $s_i^2 \leftarrow 0$, $h_i \leftarrow 0$, $i = 1, \dots, d$.
for $t = 1, \dots, T$ **do**
 Receive $\mathbf{x}_t \in \mathbb{R}^d$
 for $i = 1, \dots, d$ **do**
 $s_i^2 \leftarrow s_i^2 + x_{t,i}^2$
 if $s_i^2 > 0$ **then**
 $\eta_i \leftarrow \frac{1}{\alpha t d} \exp\left(\frac{h_i^2 + x_{t,i}^2}{2\alpha s_i^2}\right)$
 $w_i \leftarrow \eta_i \frac{h_i}{s_i^2}$
 else
 $w_i \leftarrow 0$
 Predict with $\hat{\mathbf{y}}_t = \sum_i w_i x_{t,i}$
 Receive y_t and suffer loss $\ell(y_t, \hat{\mathbf{y}}_t)$
 Compute $g_t = \partial_{\hat{\mathbf{y}}_t} \ell(y_t, \hat{\mathbf{y}}_t)$
 Update $h_i \leftarrow h_i - g_t x_{t,i}$ for all $i = 1, \dots, d$

3.2. The algorithm

We now describe an algorithm which aims at minimizing (4) for each coordinate $i = 1, \dots, d$. The algorithm maintains the negative past cumulative (linearized) losses:

$$h_{t,i} = - \sum_{q \leq t} g_q x_{q,i},$$

as well as variances $s_{t,i}^2$ for all i . At the beginning of trial t , after observing \mathbf{x}_t (and updating $s_{t,i}^2$), the algorithm predicts with weight vector \mathbf{w}_t , such that:

$$w_{t,i} = \eta_{t,i} \frac{h_{t-1,i}}{s_{t,i}^2}, \quad \text{where } \eta_{t,i} = \frac{1}{\alpha t d} \exp\left(\frac{h_{t-1,i}^2 + x_{t,i}^2}{2\alpha s_{t,i}^2}\right). \quad (5)$$

Our algorithm resembles two previously considered methods in online convex optimization, AdaptiveNormal (McMahan and Orabona, 2014) and PiSTOL (Orabona, 2014). Similarly to these methods, we also use a step size which is exponential in the square of the gradient (which is actually directly related to the same shape of regret bound (2) we are aiming for). However, we counterweight the total gradient by dividing it by the variance $s_{t,i}^2$, whereas AdaptiveNormal uses to this end the number of trials t , while PiSTOL – sum of the absolute values, $\sum_{q \leq t} |g_t x_{t,i}|$. Only our choice leads to a scale invariant algorithm, which is easiest to understand by thinking in terms of physical units: if we imagine that i -th coordinate of instances has unit $[x_i]$, the term in the exponent in (5) is unitless, while the weight w_i has unit $1/[x_i]$, so that the prediction $\hat{\mathbf{y}}_t$ also becomes unitless. Thus, rescaling the i -th coordinate (or, equivalently, changing its unit) does not affect the prediction. Note that our algorithm uses a separate “learning rate” $\eta_{t,i}$ for each coordinate, similarly to methods by McMahan and Streeter (2010); Duchi et al. (2011). The pseudo-code is presented as Algorithm 1.

We now show that the algorithm maintains small penalized regret (4). To simplify notation, define the *potential function*:

$$\psi_t(h, s) = \begin{cases} (td)^{-1} \exp\left(\frac{h^2}{2\alpha s^2}\right) & \text{when } s > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 3 *Let $\alpha_0 = \frac{9}{8}$ and define: $\kappa(\alpha) = \exp\left(\frac{1}{2(\alpha-\alpha_0)}\right)$. In each trial $t = 1, \dots, T$, for all $i = 1, \dots, d$, Algorithm 1 satisfies:*

$$g_t w_{t,i} x_{t,i} + \psi_t(h_{t,i}, s_{t,i}) \leq \psi_{t-1}(h_{t-1,i}, s_{t-1,i}) + \frac{\kappa(\alpha)}{td}.$$

The proof is given in Appendix A. Lemma 3 can be thought of as a motivation behind the particular form of the weight vector used by the algorithm: the algorithm's predictions are set to keep its loss bounded by the drop of the potential. Note, however, that the algorithm does not play with the negative gradient of the potential (which is how many online learning algorithms can be motivated), as there is additional, necessary, correction of $\exp(x_{t,i}^2/(2\alpha s_{t,i}^2))$ in the weight expression.

Applying Lemma 3 to each $t = 1, \dots, T$ and summing over trials gives:

$$\sum_t g_t w_{t,i} x_{t,i} + \psi_T(h_{T,i}, s_{T,i}) \leq \frac{\kappa(\alpha)}{d}(1 + \log T),$$

where we bound $\sum_{t=1}^T \frac{1}{t} \leq 1 + \log T$. Identifying the left-hand side of the above with the right-hand side of (4) for $\beta = Td$, and following the line of reasoning in Section 3.1, we obtain the bound on the penalized regret for the i -th coordinate:

$$\sum_t g_t x_{t,i} (w_{t,i} - u_i) \leq |u_i|_{s_{T,i}} \sqrt{\alpha \log(1 + \alpha d^2 T^2 u_i^2 s_{T,i}^2)} + \frac{\kappa(\alpha)}{d}(1 + \log T).$$

Summing over $i = 1, \dots, d$ results in the following regret bound for the algorithm:

Theorem 4 *For any comparator \mathbf{u} and any sequence of outcomes $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$, Algorithm 1 satisfies:*

$$\text{regret}_T(\mathbf{u}) \leq \sum_{i=1}^d |u_i|_{s_{T,i}} \sqrt{\alpha \log(1 + \alpha d^2 T^2 u_i^2 s_{T,i}^2)} + \kappa(\alpha)(1 + \log T).$$

We finish this section by comparing the obtained bound with a standard bound of Online Gradient Descent $UX\sqrt{T}$ when the instances and the comparator are bounded, $\|\mathbf{x}_t\| \leq X$, $\|\mathbf{u}\| \leq U$. By Cauchy-Schwarz inequality we have:

$$\sum_i |u_i|_{s_{T,i}} \leq \sqrt{\sum_i u_i^2} \sqrt{\sum_i s_{T,i}^2} \leq U \sqrt{\sum_{i,t} x_{t,i}^2} = U \sqrt{\sum_t \|\mathbf{x}_t\|^2} \leq UX\sqrt{T},$$

so that our bound is $O(UX\sqrt{T} \log(1 + d^2 U^2 X^2 T^3))$, incurring only a logarithmic overhead for not knowing the bound on the instances and on the comparator in hindsight.

4. Full Scale Invariance

In this section we consider algorithms which are invariant under general linear transformations of the form $\mathbf{x}_t \mapsto \mathbf{A}\mathbf{x}_t$ for all $t = 1, \dots, T$. As we will see, imposing such a general symmetry will lead to a second order algorithm, i.e. the algorithm will maintain the full covariance matrix $\mathbf{S}_t = \sum_{q \leq t} \mathbf{x}_q \mathbf{x}_q^\top$. To incorporate the scale invariance into the problem, we choose the penalty $\Psi_T(\mathbf{u})$ to depend only on the predictions generated by \mathbf{u} . A natural and analytically tractable choice is to parameterize $\Psi_T(\mathbf{u})$ by means of a sum of squared predictions:

$$\Psi_T(\mathbf{u}) = f(\|\mathbf{u}\|_{\mathbf{S}_T}), \quad \text{where } \|\mathbf{u}\|_{\mathbf{S}_T} \stackrel{\text{def}}{=} \sqrt{\mathbf{u}^\top \mathbf{S}_T \mathbf{u}} = \sqrt{\sum_{t=1}^T (\mathbf{x}_t^\top \mathbf{u})^2}.$$

As before, by taking into account the lower bound from Lemma 1, we choose $f(x)$ as defined in (2), i.e. $f(x) = x\sqrt{\alpha} \log(1 + \alpha\beta^2 x^2)$, for some $\alpha, \beta > 0$. Our goal is thus to design a scale-invariant algorithm which maintains small penalized regret:

$$\begin{aligned} \text{regret}_T(\mathbf{u}) - \Psi_T(\mathbf{u}) &= \sum_t g_t \mathbf{x}_t^\top \mathbf{w}_t + \mathbf{h}_T^\top \mathbf{u} - f(\|\mathbf{u}\|_{\mathbf{S}_T}) \\ &\leq \sum_t g_t \mathbf{x}_t^\top \mathbf{w}_t + \sup_{\mathbf{u}} \left(\mathbf{h}_T^\top \mathbf{u} - f(\|\mathbf{u}\|_{\mathbf{S}_T}) \right), \end{aligned}$$

where we defined $\mathbf{h}_t = -\sum_{q \leq t} g_q \mathbf{x}_q$. We will make use of the following general result, proven in the Appendix B. For any positive semi-definite matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$, let $\|\mathbf{u}\|_{\mathbf{A}} \stackrel{\text{def}}{=} \sqrt{\mathbf{u}^\top \mathbf{A} \mathbf{u}}$ denote the semi-norm of $\mathbf{u} \in \mathbb{R}^d$ induced by \mathbf{A} . We have:

Lemma 5 *For any $f(x): [0, \infty) \rightarrow \mathbb{R}$, any positive semi-definite matrix \mathbf{A} and any vector $\mathbf{y} \in \text{range}(\mathbf{A})$,*

$$\sup_{\mathbf{u}} \left\{ \mathbf{y}^\top \mathbf{u} - f(\|\mathbf{u}\|_{\mathbf{A}}) \right\} = f^*(\|\mathbf{y}\|_{\mathbf{A}^\dagger})$$

where $f^*(\theta) = \sup_{x \geq 0} x\theta - f(x)$ is the conjugate of $f(x)$ and \mathbf{A}^\dagger denotes the pseudo-inverse of \mathbf{A} . In particular,

$$\sup_{\mathbf{u}} \left\{ \mathbf{h}_T^\top \mathbf{u} - f(\|\mathbf{u}\|_{\mathbf{S}_T}) \right\} = f^*(\|\mathbf{h}_T\|_{\mathbf{S}_T^\dagger}).$$

Application of Lemma 5 together with Lemma 2 gives:

$$\text{regret}_T(\mathbf{u}) - \Psi_T(\mathbf{u}) \leq \sum_{t=1}^T g_t \mathbf{x}_t^\top \mathbf{w}_t + \underbrace{\frac{1}{\beta} \exp\left(\frac{1}{2\alpha} \|\mathbf{h}_T\|_{\mathbf{S}_T^\dagger}\right)}_{f^*(\|\mathbf{h}_T\|_{\mathbf{S}_T^\dagger})} \quad (6)$$

As before, we have eliminated the unknown comparator from the equation, and we will design the algorithm to directly minimize the right-hand side of (6) over the worst-case choice of the data.

4.1. Lower bound

We start with a negative result. It turns out that the full scale invariance setting is significantly harder than then coordinate-wise one already for $d = 2$. We will show that any algorithm will suffer at least $\Omega(\|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T})$ regret in the worst case, and this bound has a matching upper bound for a trivial algorithm which predicts 0 all the time.

Theorem 6 *Let $d \geq 2$. For any algorithm, and any nonnegative number $\beta \in \mathbb{R}_+$, there exist a sequence of outcomes and a comparator \mathbf{u} , such that $\|\mathbf{u}\|_{\mathbf{S}_T} = \beta$ and:*

$$\text{regret}_T(\mathbf{u}) \geq \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T/2}.$$

On the other hand, consider an algorithm which predicts 0 all the time. In this case,

$$\text{regret}_T(\mathbf{u}) = -\sum_{t=1}^T g_t \mathbf{x}_t^\top \mathbf{u} \leq \sum_{t=1}^T |g_t \mathbf{x}_t^\top \mathbf{u}| \leq \sqrt{\sum_t (\mathbf{x}_t^\top \mathbf{u})^2} \sqrt{\sum_t g_t^2} \leq \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T},$$

where the second inequality is from Cauchy-Schwarz inequality. Thus, the lower bound is trivially achieved, and we conclude that it is not possible to obtain meaningful bound that only depends on $\|\mathbf{u}\|_{\mathbf{S}_T}$ by any online algorithm in the worst-case.

4.2. The algorithm

While it is not possible to get a meaningful bound in terms of $\|\mathbf{u}\|_{\mathbf{S}_T}$ in the worst case, here we provide a scale-invariant algorithm which *almost* achieves that. Precisely, we derive an algorithm with a regret bound expressed by $f(\|\mathbf{u}\|_{\mathbf{S}_T})$, with f defined as in (2), with only a logarithmic dependence on the size of the instances hidden in constant β .

The algorithm is designed in order to minimize the right-hand side of (6). It maintains the negative past cumulative (linearized) loss vector $\mathbf{h}_t = -\sum_{q \leq t} g_q \mathbf{x}_q$, as well as the covariance matrix \mathbf{S}_t . Furthermore, the algorithm also keeps track of a quantity $\Gamma_t \geq 0$, recursively defined as:

$$\Gamma_0 = 0, \quad \Gamma_t = \Gamma_{t-1} + g_t^2 \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t.$$

At the beginning of trial t , after observing \mathbf{x}_t (and updating \mathbf{S}_t), the algorithm predicts with weight vector \mathbf{w}_t , such that:

$$\mathbf{w}_t = \eta_t \mathbf{S}_t^\dagger \mathbf{h}_{t-1}, \quad \text{where } \eta_t = \frac{1}{\alpha} \exp\left(\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - \Gamma_{t-1})\right). \quad (7)$$

This choice of the update leads to the invariance of the algorithm's predictions under transformations of the form $\mathbf{x}_t \mapsto \mathbf{A} \mathbf{x}_t$, $t = 1, \dots, T$, for any invertible matrix \mathbf{A} (shown in Appendix D). The algorithm is a second-order method, and is reminiscent of the Online Newton algorithm (Hazan et al., 2007; Luo et al., 2016). Our algorithm, however, adaptively chooses step size η_t ("learning rate") in each trial. Moreover, no projections are performed, which let us reduce the runtime of the algorithm to $O(d^2)$ per trial (an efficient implementation is discussed at the end of this section). The pseudo-code is presented as Algorithm 2.

Algorithm 2: Scale invariant algorithm

Parameter : $\alpha > \frac{9}{8} = 1.125$.
Initialization: $\mathbf{S} \leftarrow \mathbf{0}$, $\mathbf{h} \leftarrow \mathbf{0}$, $\Gamma \leftarrow 0$
for $t = 1, \dots, T$ **do**
 Receive $\mathbf{x}_t \in \mathbb{R}^d$
 Update $\mathbf{S} \leftarrow \mathbf{S} + \mathbf{x}_t \mathbf{x}_t^\top$
 $\eta \leftarrow \frac{1}{\alpha} \exp\left(\frac{1}{2\alpha}(\mathbf{h}^\top \mathbf{S}^\dagger \mathbf{h} - \Gamma)\right)$
 $\mathbf{w} \leftarrow \eta \mathbf{S}^\dagger \mathbf{h}$
 Predict with $\hat{y}_t = \mathbf{w}^\top \mathbf{x}_t$
 Receive y_t and suffer loss $\ell(y_t, \hat{y}_t)$
 Compute $g_t = \partial_{\hat{y}_t} \ell(y_t, \hat{y}_t)$
 Update $\mathbf{h} \leftarrow \mathbf{h} - g_t \mathbf{x}_t$
 Update $\Gamma \leftarrow \Gamma + g_t^2 \mathbf{x}_t^\top \mathbf{S}^\dagger \mathbf{x}_t$

We now bound the regret of the algorithm. Define the potential function as:

$$\psi_t(\mathbf{h}, \mathbf{S}) = \exp\left(\frac{1}{2\alpha} \mathbf{h}^\top \mathbf{S}^\dagger \mathbf{h} - \Gamma_t\right)$$

We have the following result:

Lemma 7 *In each trial $t = 1, \dots, T$, Algorithm 2 satisfies:*

$$g_t \mathbf{x}_t^\top \mathbf{w}_t + \psi_t(\mathbf{h}_t, \mathbf{S}_t) \leq \psi_{t-1}(\mathbf{h}_{t-1}, \mathbf{S}_{t-1}).$$

The proof is given in Appendix E. The choice of \mathbf{w}_t in Algorithm 2 can be motivated as the one that guarantees bounding the loss of the algorithm by the drop of the potential function (note, however, the as in the coordinate-wise case, the weight vector is not equal to the negative gradient of the potential). Comparing to Lemma 3, there is no overhead on the right-hand side; however, the overhead is actually hidden in the definition of ψ_t in quantity Γ_t .

Applying Lemma 7 to each $t = 1, \dots, T$ and summing over trials gives:

$$\sum_t g_t \mathbf{x}_t^\top \mathbf{w}_t - \psi_T(\mathbf{h}_T, \mathbf{S}_T) \leq \psi_0(\mathbf{h}_0, \mathbf{S}_0) = 1.$$

Identifying the left-hand side of the above with the right-hand side of (6) for $\beta = e^{\frac{\Gamma_T}{2\alpha}}$, we obtain the following bound on the regret:

$$\begin{aligned} \text{regret}_T(\mathbf{u}) &\leq \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{\alpha \log\left(1 + \alpha \|\mathbf{u}\|_{\mathbf{S}_T}^2 e^{\frac{\Gamma_T}{\alpha}}\right)} + 1 \\ &\leq \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{\alpha \log\left(1 + \alpha \|\mathbf{u}\|_{\mathbf{S}_T}^2\right) + \log(\alpha) \Gamma_T} + 1, \end{aligned}$$

where we used $\log(1+ab) \leq \log(a+ab) = \log a + \log(1+b)$ for $a \geq 1$, applied to $a = e^{\Gamma_T/\alpha} \geq 1$. Thus, the algorithm achieves an essentially optimal (up to logarithmic factor) scale-invariant bound expressed in terms $\|\mathbf{u}\|_{\mathbf{S}_T}$, with an additional overhead hidden in Γ_T .

How large can Γ_T be? By the definition, $\Gamma_T = \sum_t g_t^2 \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t$; as $g_t^2 \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t \leq g_t^2 \leq 1$, Γ_T is at most T in the worst case, and the bound becomes $O(\|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{T})$ (logarithmic factors dropped), which is what we expected given the negative result in Theorem 6. However, Γ_T can be much smaller in most practical cases as it can be shown to grow only *logarithmically* with the size of the instances (Luo et al., 2016, notation translated to our setup):

Lemma 8 (Luo et al., 2016, Theorem 4) *Let λ^* be the minimum among the smallest nonzero eigenvalues of \mathbf{S}_t ($t = 1, \dots, T$) and r be the rank of \mathbf{S}_T . We have:*

$$\sum_{t=1}^T \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t \leq r + \frac{(1+r)r}{2} \log \left(1 + \frac{2 \sum_{t=1}^T \|\mathbf{x}_t\|^2}{(1+r)r\lambda^*} \right).$$

Combining the above results, we thus get:

Theorem 9 *For any comparator \mathbf{u} and any sequence of outcomes $\{(\mathbf{x}_t, y_t)\}_{t=1}^T$, Algorithm 2 satisfies:*

$$\text{regret}_T(\mathbf{u}) \leq \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{\alpha \log(1 + \alpha \|\mathbf{u}\|_{\mathbf{S}_T}^2) + \log(\alpha) \Gamma_T} + 1,$$

where:

$$\Gamma_T = \sum_{t=1}^T g_t \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t \leq r + \frac{(1+r)r}{2} \log \left(1 + \frac{2 \sum_{t=1}^T \|\mathbf{x}_t\|^2}{(1+r)r\lambda^*} \right).$$

with λ^* being the minimum among the smallest nonzero eigenvalues of \mathbf{S}_t ($t = 1, \dots, T$) and r being the rank of \mathbf{S}_T .

We finally note that the dependence on the dimension d in the bound (through the dependence on the rank r in Γ_T) cannot be eliminated, as Luo et al. (2016, Theorem 1) show that in a setting in which the predictions of \mathbf{u} are constrained to be at most C , any algorithm will suffer the regret at least $\Omega(C\sqrt{dT})$.⁴

Efficient implementation. The dominating cost in Algorithm 2 is the computation of pseudoinverse \mathbf{S}_t^\dagger in each trial after performing the update $\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$, which can be $O(d^3)$. However, we can improve the computational cost per trial to $O(d^2)$ by noticing that \mathbf{S}_t is never used by the algorithm, so it suffices to store and directly update \mathbf{S}_t^\dagger using a rank-one update procedure in the spirit Sherman-Morrison formula, which takes $O(d^2)$. The procedure is highlighted in the proof of Lemma 7 in Appendix E.

5. Conclusions

We considered unconstrained online convex optimization, exploiting a natural scale invariance symmetry: the predictions of the optimal comparator (weight vector) are invariant under any linear transformation of the instances (input vectors). Thus, the scale of the

4. We can, however, improve the dependence on d to $O(\sqrt{d})$ by modifying the algorithm to play with $\tilde{\mathbf{S}}_t = \epsilon \mathbf{I} + \mathbf{S}_t$ for $\epsilon > 0$, and apply the bound on $\sum_t \mathbf{x}_t^\top \tilde{\mathbf{S}}_t^\dagger \mathbf{x}_t$ from Cesa-Bianchi and Lugosi (2006), Theorem 11.7. This would, however, come at the price of losing the scale invariance of the algorithm.

weight vector is only relative to the scale of the instances, and we aimed at designing online algorithms which also enjoy this property, i.e. are scale-invariant. We first considered the case of coordinate-wise invariance, in which the individual coordinates (features) can be arbitrarily rescaled. We gave an algorithm, which achieves essentially (up to logarithmic factor) optimal regret bound in this setup (expressed by means of a coordinate-wise scale-invariant norm of the comparator). We then moved to a general (full) invariance with respect to arbitrary linear transformations. We first gave a negative result, showing that no algorithm can achieve a meaningful bound in terms of scale-invariant norm of the comparator in the worst case. Next, we complimented this result with a positive one, providing an algorithm which “almost” achieve the desired bound, incurring only a logarithmic overhead in terms of the norm of the instances.

In the future research, we plan to test the introduced algorithms in the computational experiments to verify how their performance relate to the existing online methods from the past work (Zinkevich, 2003; Ross et al., 2013; Orabona et al., 2015; Luo et al., 2016).

Acknowledgments

We thank the anonymous reviewers for suggestions which improved the quality of our work. The author acknowledges support from the Polish National Science Centre (grant no. 2016/22/E/ST6/00299).

References

- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- Stephen L. Campbell and Carl D. Meyer. *Generalized Inverses of Linear Transformations*. SIAM, 2009.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- Xuzhou Chen and Jun Ji. Computing the Moore-Penrose inverse of a matrix through symmetric rank-one updates. *American Journal of Computational Mathematics*, 1(3):147–151, 2011.
- John C. Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends in Optimization*, 2(3–4):157–325, 2015.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.
- Jyrki Kivinen and Manfred K. Warmuth. Additive versus Exponentiated Gradient updates for linear prediction. *Information and Computation*, 132(1):1–64, 1997.

- Haipeng Luo, Alekh Agarwal, Nicolás Cesa-Bianchi, and John Langford. Efficient second order online learning by sketching. In *Advances in Neural Information Processing Systems (NIPS) 29*, 2016.
- H. Brendan McMahan and Jacob Abernethy. Minimax optimal algorithms for unconstrained linear optimization. In *Advances in Neural Information Processing Systems (NIPS) 26*, pages 2724–2732, 2013.
- H. Brendan McMahan and Francesco Orabona. Unconstrained online linear learning in Hilbert spaces: Minimax algorithms and normal approximation. In *Proc. of the 27th Conference on Learning Theory (COLT)*, pages 1020–1039, 2014.
- H. Brendan McMahan and Matthew J. Streeter. Adaptive bound optimization for online convex optimization. In *Conference on Learning Theory (COLT)*, pages 244–256, 2010.
- Francesco Orabona. Dimension-free exponentiated gradient. In *Advances in Neural Information Processing Systems (NIPS) 26*, pages 1806–1814, 2013.
- Francesco Orabona. Simultaneous model selection and optimization through parameter-free stochastic learning. In *Advances in Neural Information Processing Systems (NIPS) 27*, pages 1116–1124, 2014.
- Francesco Orabona and Dávid Pál. Scale-free algorithms for online linear optimization. In *Algorithmic Learning Theory (ALT)*, pages 287–301, 2015.
- Francesco Orabona and Dávid Pál. Coin betting and parameter-free online learning. In *Neural Information Processing Systems (NIPS)*, 2016.
- Francesco Orabona, Koby Crammer, and Nicolò Cesa-Bianchi. A generalized online mirror descent with applications to classification and regression. *Machine Learning*, 99(3):411–435, 2015.
- Stephane Ross, Paul Mineiro, and John Langford. Normalized online learning. In *Proc. of the 29th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 537–545, 2013.
- Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2011.
- Matthew Streeter and H. Brendan McMahan. No-regret algorithms for unconstrained online convex optimization. In *Advances in Neural Information Processing Systems (NIPS) 25*, pages 2402–2410, 2012.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *International Conference on Machine Learning (ICML)*, pages 928–936, 2003.

Appendix A. Proof of Lemma 3

Let t_0 be the first trial in which $s_{t_0,i} > 0$. This means that in trials $t = 1, \dots, t_0 - 1$, $x_{t,i} = 0$ and hence $h_{t,i} = 0$, and the lemma is trivially satisfied, as the left-hand side is zero. At trial t_0 , the algorithm still predicts with $w_{t,i} = 0$, and Lemma 3 boils down to showing that

$$\frac{1}{t_0 d} \exp\left(\frac{h_{t_0,i}^2}{2\alpha s_{t_0,i}^2}\right) \leq \frac{\kappa(\alpha)}{t_0 d},$$

where $\kappa(\alpha) = e^{\frac{1}{2(\alpha-\alpha_0)}}$. Since $h_{t_0,i} = -g_{t_0} x_{t_0,i}$ and $s_{t_0,i}^2 = x_{t_0,i}^2$, we have $\frac{h_{t_0,i}^2}{s_{t_0,i}^2} \leq 1$, and thus the left-hand side is bounded by:

$$\frac{1}{t_0 d} e^{\frac{1}{2\alpha}} \leq \frac{1}{t_0 d} e^{\frac{1}{2(\alpha-\alpha_0)}} = \frac{\kappa(\alpha)}{t_0 d}.$$

Thus, we have shown the lemma for trials $t = 1, \dots, t_0$. We can now assume that $s_{t-1,i} > 0$ and prove the lemma for the remaining trials $t > t_0$. By using the definition of $w_{t,i}$ from (5), we need to show:

$$\frac{g_t x_{t,i} h_{t-1,i}}{s_{t,i}^2 \alpha t d} \exp\left(\frac{h_{t-1,i}^2 + x_{t,i}^2}{2\alpha s_{t,i}^2}\right) + \frac{1}{t d} \exp\left(\frac{h_{t,i}^2}{2\alpha s_{t,i}^2}\right) \leq \frac{1}{(t-1)d} \exp\left(\frac{h_{t-1,i}^2}{2\alpha s_{t-1,i}^2}\right) + \frac{\kappa(\alpha)}{t d}, \quad (8)$$

where we remind that $h_{t,i} = h_{t-1,i} - g_t x_{t,i}$ and $s_{t,i}^2 = s_{t-1,i}^2 + x_{t,i}^2$. First note that the left-hand side is convex in g_t , and hence it is maximized for $g_t \in \{-1, 1\}$. As the right-hand side does not depend on g_t , it suffices to show that the inequality holds for $g_t \in \{-1, 1\}$. Furthermore, as the inequality depends on g_t only through the product $g_t x_{t,i}$, we assume without loss of generality that $x_{t,i} \geq 0$ (the sign can always be incorporated to g_t). We now simplify the notation. Define:

$$v = \frac{g_t h_{t-1,i}}{s_{t-1,i}}, \quad \gamma^2 = \frac{s_{t-1,i}^2}{s_{t,i}^2}.$$

Note that by the definition $\gamma \in (0, 1]$ as $x_{t,i}$ is unconstrained. In this notation, we have:

$$\begin{aligned} \frac{x_{t,i}}{s_{t,i}} &= \sqrt{\frac{x_{t,i}^2}{s_{t,i}^2}} = \sqrt{1 - \frac{s_{t-1,i}^2}{s_{t,i}^2}} = \sqrt{1 - \gamma^2}, \\ \frac{h_{t-1,i}^2}{s_{t,i}^2} &= \frac{h_{t-1,i}^2}{s_{t-1,i}^2} \frac{s_{t-1,i}^2}{s_{t,i}^2} = v^2 \gamma^2 \\ \frac{h_{t,i}^2}{s_{t,i}^2} &= \frac{h_{t-1,i}^2}{s_{t,i}^2} - 2 \frac{h_{t-1,i} g_t x_{t,i}}{s_{t,i}^2} + \frac{x_{t,i}^2}{s_{t,i}^2} = v^2 \gamma^2 - 2v\gamma\sqrt{1-\gamma^2} + (1-\gamma^2), \end{aligned}$$

where we used $g_t^2 = 1$. Using the new notation in (8), multiplying both sides by td we equivalently get:

$$\frac{v\gamma\sqrt{1-\gamma^2}}{\alpha} e^{\frac{v^2\gamma^2+1-\gamma^2}{2\alpha}} + e^{\frac{v^2\gamma^2-2v\gamma\sqrt{1-\gamma^2}+1-\gamma^2}{2\alpha}} \leq \frac{t}{t-1} e^{\frac{v^2}{2\alpha}} + \kappa(\alpha). \quad (9)$$

Let us denote the left-hand side of (9) as A . We have:

$$A = e^{\frac{v^2\gamma^2+1-\gamma^2}{2\alpha}} \left(\frac{v\gamma\sqrt{1-\gamma^2}}{\alpha} + e^{\frac{-v\gamma\sqrt{1-\gamma^2}}{\alpha}} \right)$$

We need the following result, which is proved in Appendix F:

Lemma 10 *Let $\alpha_0 = \frac{9}{8}$. For all $x \in \mathbb{R}$ it holds:*

$$x + e^{-x} \leq e^{\frac{x^2}{2}\alpha_0}.$$

Using Lemma 10 with $x = \frac{v\gamma\sqrt{1-\gamma^2}}{\alpha}$, we bound:

$$A \leq \exp \left\{ \frac{1}{2\alpha} \underbrace{(v^2\gamma^2 + 1 - \gamma^2 + av^2\gamma^2(1 - \gamma^2))}_{=B(\gamma^2)} \right\},$$

where $a = \frac{\alpha_0}{\alpha} < 1$. Note that $B(\gamma^2)$ is a concave quadratic function of γ^2 , and its (unconstrained) maximizer is given by: $\gamma^{*2} = \frac{1}{2} + \frac{v^2-1}{2av^2}$. However, since the allowed values of γ^2 are $(0, 1]$, the maximizer within the *closure* of this range is:

- $\gamma^{*2} = 0$ for $v^2 \leq \frac{1}{1+a}$, which gives $B(\gamma^{*2}) = 1$,
- $\gamma^{*2} = 1$ for $v^2 \geq \frac{1}{1-a}$, which gives $B(\gamma^{*2}) = v^2$,
- $\gamma^{*2} = \frac{1}{2} + \frac{v^2-1}{2av^2}$ for $v^2 \in [\frac{1}{1+a}, \frac{1}{1-a}]$, which gives $B(\gamma^{*2}) = \frac{a-1}{2a} + v^2 \frac{(1+a)^2}{4a} + \frac{1}{4av^2}$. Since this function is monotonically increasing in v^2 for $v^2 \geq \frac{1}{1+a}$, we will upper bound it by setting $v^2 = \frac{1}{1-a}$, which gives $B(\gamma^*) \leq \frac{1}{1-a}$.

We thus jointly upper bound $B(\gamma) \leq \max\{v^2, \frac{1}{1-a}\}$, which results in:

$$A \leq \max \left\{ e^{\frac{v^2}{2\alpha}}, e^{\frac{1}{2\alpha(1-a)}} \right\} \leq e^{\frac{v^2}{2\alpha}} + e^{\frac{1}{2\alpha(1-a)}} \leq \frac{t}{t-1} e^{\frac{v^2}{2\alpha}} + \kappa(\alpha),$$

which verifies (9) and finishes the proof.

Appendix B. Proof of Lemma 5

Without loss of generality assume \mathbf{A} has $k \geq 1$ strictly positive eigenvalues (the remaining eigenvalues being zero), as otherwise (for $k = 0$) $\text{range}(\mathbf{A}) = \{\mathbf{0}\}$ and there is nothing to show. Let $\mathbf{A} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$ for $\mathbf{V} \in \mathbb{R}^{d \times k}$, $\mathbf{\Sigma} = \text{diag}(\lambda_1, \dots, \lambda_k)$, be the ‘thin’ eigenvalue decomposition of \mathbf{A} (i.e., the eigendecomposition without explicit appearance of eigenvectors with zero eigenvalues). Since $\mathbf{y} \in \text{range}(\mathbf{A})$, there exists $\tilde{\mathbf{y}} \in \mathbb{R}^d$, such that $\mathbf{y} = \mathbf{A}\tilde{\mathbf{y}}$, and therefore:

$$\mathbf{y} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top\tilde{\mathbf{y}} = \mathbf{V}\mathbf{\Sigma}^{1/2}\mathbf{z}, \quad \text{where } \mathbf{z} = \mathbf{\Sigma}^{1/2}\mathbf{V}^\top\tilde{\mathbf{y}} \in \mathbb{R}^k.$$

We have:

$$\begin{aligned} \sup_{\mathbf{u} \in \mathbb{R}^d} \left\{ \mathbf{y}^\top \mathbf{u} - f\left(\sqrt{\mathbf{u}^\top \mathbf{A} \mathbf{u}}\right) \right\} &= \sup_{\mathbf{u} \in \mathbb{R}^d} \left\{ \mathbf{z}^\top (\boldsymbol{\Sigma}^{1/2} \mathbf{V}^\top \mathbf{u}) - f\left(\|\boldsymbol{\Sigma}^{1/2} \mathbf{V}^\top \mathbf{u}\|\right) \right\} \\ &= \sup_{\mathbf{u} \in \mathbb{R}^k} \left\{ \mathbf{z}^\top \mathbf{u} - f(\|\mathbf{u}\|) \right\}, \end{aligned}$$

where we reparametrized $\boldsymbol{\Sigma}^{1/2} \mathbf{V}^\top \mathbf{u}$ as $\mathbf{u} \in \mathbb{R}^k$. Now, note that keeping the norm $\|\mathbf{u}\|$ fixed, the supremum is achieved by \mathbf{u} in the direction of \mathbf{z} ; therefore, without loss of generality assume $\mathbf{u} = \beta \frac{\mathbf{z}}{\|\mathbf{z}\|}$ for some $\beta \geq 0$. This means that:

$$\sup_{\mathbf{u} \in \mathbb{R}^d} \left\{ \mathbf{z}^\top \mathbf{u} - f(\|\mathbf{u}\|) \right\} = \sup_{\beta \geq 0} \left\{ \beta \|\mathbf{z}\| - f(\beta) \right\} = f^*(\|\mathbf{z}\|).$$

Since $\mathbf{A}^\dagger = \mathbf{V} \boldsymbol{\Sigma}^{-1} \mathbf{V}^\top$,

$$\|\mathbf{y}\|_{\mathbf{A}^\dagger}^2 = \mathbf{y}^\top \mathbf{A}^\dagger \mathbf{y} = \mathbf{y}^\top \mathbf{V} \boldsymbol{\Sigma}^{-1} \mathbf{V}^\top \mathbf{y} = \|\boldsymbol{\Sigma}^{-1/2} \mathbf{V}^\top \mathbf{y}\|^2 = \|\mathbf{z}\|^2,$$

which finishes the proof of the first part of the lemma.

For the second part, we only need $\mathbf{h}_T \in \text{range}(\mathbf{S}_T)$, a well-known fact, which we show below for completeness. Let \mathbf{v} be any eigenvector of \mathbf{S}_T associated with zero eigenvalue, so that $\mathbf{v}^\top \mathbf{S}_T \mathbf{v} = 0$. Using the definition $\mathbf{S}_T = \sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top$, we have $\mathbf{v}^\top \mathbf{S}_T \mathbf{v} = \sum_{t=1}^T (\mathbf{x}_t^\top \mathbf{v})^2 = 0$, which implies $\mathbf{x}_t^\top \mathbf{v} = 0$ for all t . But since $\mathbf{h}_T = -\sum_{t=1}^T g_t \mathbf{x}_t$, this also means $\mathbf{h}_T^\top \mathbf{v} = 0$. Let $\mathbf{v}_1, \dots, \mathbf{v}_k$ be the eigenvectors of \mathbf{S}_T associated with all non-zero eigenvalues $\lambda_1, \dots, \lambda_k$. By the previous argument, $\mathbf{h}_T \in \text{span}\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$, i.e. $\mathbf{h}_T = \sum_{i=1}^k \alpha_i \mathbf{v}_i$. Choosing a vector $\mathbf{z} = \sum_{i=1}^k \frac{\alpha_i}{\lambda_i} \mathbf{v}_i$ reveals that:

$$\mathbf{S}_T \mathbf{z} = \sum_{i=1}^k \lambda_i \mathbf{v}_i \mathbf{v}_i^\top \mathbf{z} = \sum_{i=1}^k \lambda_i \frac{\alpha_i}{\lambda_i} \mathbf{v}_i = \sum_{i=1}^k \alpha_i \mathbf{v}_i = \mathbf{h}_T,$$

which shows that $\mathbf{h}_T \in \text{range}(\mathbf{S}_T)$. This finishes the proof.

Appendix C. Proof of Theorem 6

We will show that for any algorithm, there exists a sequence of outcomes $\{(\mathbf{x}_t, g_t)\}_{t=1}^T$, such that the loss of the algorithm is nonnegative on this sequence, while $\mathbf{h}_T^\top \mathbf{S}_T^\dagger \mathbf{h}_T = \frac{T}{2}$. Then, we can choose the comparator as $\mathbf{u} = \beta \sqrt{\frac{2}{T}} \mathbf{S}_T^\dagger \mathbf{h}_T$, which invariant norm is equal to:

$$\|\mathbf{u}\|_{\mathbf{S}_T} = \sqrt{\mathbf{u}^\top \mathbf{S}_T \mathbf{u}} = \beta \sqrt{\frac{2}{T}} \sqrt{\mathbf{h}_T^\top \mathbf{S}_T^\dagger \mathbf{S}_T \mathbf{S}_T^\dagger \mathbf{h}_T} = \beta \sqrt{\frac{2}{T}} \sqrt{\mathbf{h}_T^\top \mathbf{S}_T^\dagger \mathbf{h}_T} = \beta \sqrt{\frac{2}{T}} \sqrt{\frac{T}{2}} = \beta,$$

and the regret becomes:

$$\text{regret}_T(\mathbf{u}) = \underbrace{\sum_{t=1}^T g_t \mathbf{x}_t^\top \mathbf{w}_t}_{\text{nonnegative}} + \mathbf{h}_T^\top \mathbf{u} \geq \mathbf{h}_T^\top \mathbf{u} = \beta \sqrt{\frac{2}{T}} \mathbf{h}_T^\top \mathbf{S}_T^\dagger \mathbf{h}_T = \beta \sqrt{\frac{T}{2}} = \|\mathbf{u}\|_{\mathbf{S}_T} \sqrt{\frac{T}{2}}.$$

Thus, to finish the proof, it suffices to find a sequence with the claimed properties. The sequence goes as follows. In each trial t , the adversary chooses the sign of g_t to match the sign of $\hat{y}_t = \mathbf{x}_t^\top \mathbf{w}_t$, so that $g_t \mathbf{x}_t^\top \mathbf{w}_t \geq 0$, i.e. the loss of the algorithm is nonnegative in each trial. Moreover, \mathbf{x}_t and g_t are chosen so that $\mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t = \frac{t}{2}$ for all t , and hence $\mathbf{h}_T^\top \mathbf{S}_T^\dagger \mathbf{h}_T = \frac{T}{2}$. In the first d trials, the adversary chooses $\mathbf{x}_t = \mathbf{e}_t$ and $g_t \in \{-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}\}$. In this case, for any $t = 1, \dots, d$:

$$\mathbf{h}_t = (g_1, \dots, g_t, 0, \dots, 0), \quad \mathbf{S}_t = \mathbf{S}_t^\dagger = \text{diag}(\underbrace{1, \dots, 1}_t, 0, \dots, 0),$$

and hence $\mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t = \sum_{i=1}^t g_i^2 = \frac{t}{2}$. Note that \mathbf{G}_t is invertible for all $t \geq d$, i.e. $\mathbf{S}_t^\dagger = \mathbf{S}_t^{-1}$ for $t \geq d$. For the remaining trials $t = d+1, \dots, T$ the adversary chooses $g_t \in \{-1, 1\}$, and \mathbf{x}_t in any direction for which $\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t = 0$, with $\|\mathbf{x}_t\|$ chosen in such a way that $\mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t = 1$ (which is always possible as $d \geq 2$). Using Sherman-Morrison inversion formula for $\mathbf{S}_t = \mathbf{S}_{t-1} + \mathbf{x}_t \mathbf{x}_t^\top$,

$$\begin{aligned} \mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t &= \mathbf{h}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{h}_t - \frac{(\mathbf{h}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t)^2}{1 + \mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t} = \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^{-1} \mathbf{h}_{t-1} + \mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t - \frac{(\mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t)^2}{1 + \mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t} \\ &= \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^{-1} \mathbf{h}_{t-1} + \frac{\mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t}{1 + \mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t} = \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1} + \frac{1}{2}, \end{aligned}$$

where in the second equality we used $\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t = 0$ and in the last equality we used $\mathbf{x}_t^\top \mathbf{S}_{t-1}^{-1} \mathbf{x}_t = 1$. Thus, $\mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t = \frac{t}{2}$ for all t .

Appendix D. Scale invariance of Algorithm 2

We need to show that under linear transformation $\mathbf{x}_t \mapsto \mathbf{A} \mathbf{x}_t$, $t = 1, \dots, T$, for any invertible \mathbf{A} , the predictions of the algorithm, $\{\hat{y}_t\}_{t=1}^T$, do not change. We remind that the predictions are given by:

$$\hat{y}_t = \eta_t \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}, \quad \text{where } \eta_t = \frac{1}{\alpha} \exp\left(\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - \Gamma_{t-1})\right), \quad \Gamma_t = \sum_{q \leq t} g_q^2 \mathbf{x}_q^\top \mathbf{S}_t^\dagger \mathbf{x}_q.$$

We proceed by induction on t . For $t = 1$, $\hat{y}_1 = 0$, which is trivially invariant under linear transformation of the instances. Now, assume inductively that $\{\hat{y}_q\}_{q=1}^{t-1}$ are invariant, which implies $\{g_q\}_{q=1}^{t-1}$ are also invariant as they only depend on past predictions and past labels. Since $\mathbf{h}_{t-1} = -\sum_{q < t} g_q \mathbf{x}_q$, prediction \hat{y}_t depends on the data only by means of $\{g_q\}_{q=1}^{t-1}$ and quantities $\mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{x}_j$ for $i, j \leq t$. Thus, to show the invariance of \hat{y}_t under linear transformation of the instances, it suffices to show the invariance of $\mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{x}_j$, for $i, j \leq t$. As $\mathbf{S}_t = \sum_{q \leq t} \mathbf{x}_q \mathbf{x}_q^\top$ maps to $\sum_{q \leq t} \mathbf{A} \mathbf{x}_q \mathbf{x}_q^\top \mathbf{A}^\top = \mathbf{A} \mathbf{S}_t \mathbf{A}^\top$ under linear transformation, it thus amounts to show that:

$$\mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{x}_j = (\mathbf{A} \mathbf{x}_i)^\top (\mathbf{A} \mathbf{S}_t \mathbf{A}^\top)^\dagger \mathbf{A} \mathbf{x}_j = \mathbf{x}_i^\top \mathbf{A}^\top (\mathbf{A} \mathbf{S}_t \mathbf{A}^\top)^\dagger \mathbf{A} \mathbf{x}_j,$$

for any $i, j \leq t$ and any invertible \mathbf{A} .⁵ Since $\mathbf{x}_i, \mathbf{x}_j \in \text{range}(\mathbf{S}_t)$, we have: $\mathbf{x}_i = \mathbf{S}_t^\dagger \mathbf{S}_t \mathbf{x}_i$ and similarly $\mathbf{x}_j = \mathbf{S}_t^\dagger \mathbf{S}_t \mathbf{x}_j$. Thus:

$$\begin{aligned} \mathbf{x}_i^\top \mathbf{A}^\top \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right)^\dagger \mathbf{A} \mathbf{x}_j &= \mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{S}_t \mathbf{A}^\top \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right)^\dagger \mathbf{A} \mathbf{S}_t \mathbf{S}_t^\dagger \mathbf{x}_j \\ &= \mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{A}^{-1} \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right) \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right)^\dagger \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right) \mathbf{A}^{-\top} \mathbf{S}_t^\dagger \mathbf{x}_j \\ &= \mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{A}^{-1} \left(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top \right) \mathbf{A}^{-\top} \mathbf{S}_t^\dagger \mathbf{x}_j \\ &= \mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{S}_t \mathbf{S}_t^\dagger \mathbf{x}_j \\ &= \mathbf{x}_i^\top \mathbf{S}_t^\dagger \mathbf{x}_j, \end{aligned}$$

which was to be shown.

Appendix E. Proof of Lemma 7

By plugging the definition of the algorithm's weight (7) to the inequality in the lemma, we need to show:

$$\frac{g_t \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}}{\alpha} e^{\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - \Gamma_{t-1})} + e^{\frac{1}{2\alpha} (\mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t - \Gamma_t)} \leq e^{\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - \Gamma_{t-1})}. \quad (10)$$

Using $\mathbf{h}_t = \mathbf{h}_{t-1} - g_t \mathbf{x}_t$, we have:

$$\begin{aligned} \mathbf{h}_t^\top \mathbf{S}_t^\dagger \mathbf{h}_t - \Gamma_t &= \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - 2g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t + g_t^2 \mathbf{x}_t^\top \mathbf{S}_t^\dagger \mathbf{x}_t - \Gamma_t \\ &= \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} - 2g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t - \Gamma_{t-1}, \end{aligned}$$

from the definition of Γ_t . Plugging the above into (10) and multiplying both sides by $e^{\Gamma_{t-1}}$, we equivalently need to show:

$$e^{\frac{1}{2\alpha} \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}} \left(\frac{g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t}{\alpha} + e^{-\frac{g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t}{\alpha}} \right) \leq e^{\frac{1}{2\alpha} \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}}. \quad (11)$$

Denote the left-hand side of (11) by A . Applying $x + e^{-x} \leq e^{\frac{x^2}{2} \alpha_0}$ for $\alpha_0 = \frac{9}{8}$ (Lemma 10) with $x = \frac{g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t}{\alpha}$ to the left-hand side of (11) results in the following bound:

$$\begin{aligned} A &\leq \exp \left(\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}) + \frac{\alpha_0}{\alpha} (g_t \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t)^2 \right) \\ &\leq \exp \left(\frac{1}{2\alpha} (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1}) + (\mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t)^2 \right), \end{aligned}$$

where we used $\alpha_0 \leq \alpha$ and $g_t^2 \leq 1$.

We now express \mathbf{S}_t^\dagger in terms of \mathbf{S}_{t-1}^\dagger , by using an extension of well-known Sherman-Morrisson formula to pseudoinverse (Campbell and Meyer, 2009; Chen and Ji, 2011). To

5. The simplest approach to show the invariance would be to prove $(\mathbf{A} \mathbf{S}_t \mathbf{A}^\top)^\dagger = \mathbf{A}^{-\top} \mathbf{S}_t^\dagger \mathbf{A}^{-1}$. Unfortunately (and surprisingly to us), while this relation holds for a matrix inverse, it *does not* hold in general for pseudoinverse! In the proof, we need to use a crucial fact that \mathbf{x}_i and \mathbf{x}_j are in the range of \mathbf{S}_t .

this end define $\mathbf{x}_\perp = (\mathbf{I} - \mathbf{S}_{t-1}\mathbf{S}_{t-1}^\dagger)\mathbf{x}_t$ to be the component of \mathbf{x}_t orthogonal to the range of \mathbf{S}_{t-1} . Note that since $\mathbf{h}_{t-1} \in \text{range}(\mathbf{S}_{t-1})$ (see the proof of Lemma 5 in Appendix B), we have $\mathbf{h}_{t-1}^\top \mathbf{x}_\perp = 0$. Moreover, $\mathbf{x}_\perp^\top \mathbf{x}_t = \mathbf{x}_t^\top (\mathbf{I} - \mathbf{S}_{t-1}\mathbf{S}_{t-1}^\dagger)\mathbf{x}_t = \mathbf{x}_t^\top (\mathbf{I} - \mathbf{S}_{t-1}\mathbf{S}_{t-1}^\dagger)^2 \mathbf{x}_t = \mathbf{x}_\perp^\top \mathbf{x}_\perp = \|\mathbf{x}_\perp\|^2$, where we used that $\mathbf{I} - \mathbf{S}_{t-1}\mathbf{S}_{t-1}^\dagger = (\mathbf{I} - \mathbf{S}_{t-1}\mathbf{S}_{t-1}^\dagger)^2$ is idempotent as a projection operator. Let us also define $\beta = 1 + \mathbf{x}_t^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t$. Depending on whether $\mathbf{x}_\perp = \mathbf{0}$ we need to consider two cases (Chen and Ji, 2011):

1. Case $\mathbf{x}_\perp = \mathbf{0}$. In this case, we essentially follow Sherman-Morrison formula:

$$\mathbf{S}_t^\dagger = \mathbf{S}_{t-1}^\dagger - \frac{1}{\beta} \mathbf{S}_{t-1}^\dagger \mathbf{x}_t \mathbf{x}_t^\top \mathbf{S}_{t-1}^\dagger,$$

and we get:

$$\begin{aligned} \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} &= \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1} - \frac{1}{\beta} (\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t)^2, \\ \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t &= \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t - \frac{1}{\beta} \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t \underbrace{\mathbf{x}_t^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t}_{=\beta-1} = \frac{1}{\beta} \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t. \end{aligned}$$

Therefore:

$$\begin{aligned} A &\leq \exp\left(\frac{1}{2\alpha} \left(\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1} - \frac{1}{\beta} (\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t)^2 + \frac{1}{\beta^2} (\mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t)^2 \right)\right) \\ &\leq \exp\left(\frac{1}{2\alpha} \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1}\right), \end{aligned}$$

where we used $\beta \geq 1$. Thus, (11) follows.

2. Case $\mathbf{x}_\perp \neq \mathbf{0}$. In this case the update formula has a different form:

$$\mathbf{S}_t^\dagger = \mathbf{S}_{t-1}^\dagger - \frac{\mathbf{S}_{t-1}^\dagger \mathbf{x}_t \mathbf{x}_\perp^\top}{\|\mathbf{x}_\perp\|^2} - \frac{\mathbf{x}_\perp \mathbf{x}_t^\top \mathbf{S}_{t-1}^\dagger}{\|\mathbf{x}_\perp\|^2} + \beta \frac{\mathbf{x}_\perp \mathbf{x}_\perp^\top}{\|\mathbf{x}_\perp\|^4},$$

and we get:

$$\begin{aligned} \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{h}_{t-1} &= \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1}, \\ \mathbf{h}_{t-1}^\top \mathbf{S}_t^\dagger \mathbf{x}_t &= \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t - \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{x}_t \frac{\mathbf{x}_\perp^\top \mathbf{x}_t}{\|\mathbf{x}_\perp\|^2} = 0, \end{aligned}$$

where we used $\mathbf{h}^\top \mathbf{x}_\perp = 0$ and $\mathbf{x}_\perp^\top \mathbf{x}_t = \|\mathbf{x}_\perp\|^2$. Therefore,

$$A \leq \exp\left(\frac{1}{2\alpha} \mathbf{h}_{t-1}^\top \mathbf{S}_{t-1}^\dagger \mathbf{h}_{t-1}\right),$$

and (11) follows. This finishes the proof.

Appendix F. Proof of Lemma 10

We need to show that for all $x \in \mathbb{R}$ it holds:

$$x + e^{-x} \leq e^{\frac{x^2}{2}\alpha_0}. \quad (12)$$

where $\alpha_0 = \frac{9}{8}$. First, consider $x > 0$. We have $e^{\frac{x^2}{2}\alpha_0} \geq e^{\frac{x^2}{2}} \geq 1 + \frac{x^2}{2}$, and $e^x \geq 1 + x + \frac{x^2}{2}$, which implies $e^{-x} \leq \frac{1}{1+x+\frac{x^2}{2}}$. Thus, it suffices to show $1 + \frac{x^2}{2} \geq \frac{1}{1+x+\frac{x^2}{2}} + x$, which amounts to:

$$\frac{x^2}{2} - x + 1 \geq \frac{1}{1+x+\frac{x^2}{2}} \iff \left(\frac{x^2}{2} + 1 - x\right)\left(\frac{x^2}{2} + 1 + x\right) \geq 1 \iff \frac{x^4}{4} \geq 0,$$

which clearly holds. Now, consider $x \leq 0$. Showing (12) for $x < 0$ is equivalent to showing:

$$e^{\frac{x^2}{2}\alpha_0} + x - e^x \geq 0, \quad (13)$$

for $x \geq 0$ (after substituting $x \rightarrow -x$). First note that when $x \geq \frac{2}{\alpha_0}$, $e^{\frac{x^2}{2}\alpha_0} \geq e^x \geq e^x - x$, so that (13) holds. Thus, it suffices to check the inequality for $x \in [0, 2/\alpha_0] = [0, 16/9]$. When $x \leq 0.34$, we make use of the fact that $\frac{e^x - x - 1}{x^2}$ is nondecreasing in x , so that $e^x - x \leq 1 + x^2 \frac{e^{0.34} - 0.34 - 1}{0.34^2} \leq 1 + 0.5619x^2$, whereas $e^{\frac{x^2}{2}\alpha_0} \geq 1 + \frac{x^2}{2}\alpha_0 \geq 1 + \frac{9}{16}x^2 = 1 + 0.5625x^2$. Therefore, $e^{\frac{x^2}{2}\alpha_0} - e^x + x \geq 0.0006x^2 \geq 0$, and we showed (13) for $x \leq 0.34$. From now on, the proof becomes very tedious and requires first-order Taylor approximations and the use of convexity at various subintervals of $[0.34, \frac{16}{9}]$ to finally combine the bound. Thus, the rest of the proof works by splitting the range $[0.34, \frac{16}{9}]$ into intervals $[u_1, v_1], \dots, [u_m, v_m]$, where $u_1 = 0.34$, $u_i = v_{i-1}$, and $v_m = \frac{16}{9}$. For each $i = 1, \dots, m$, for $x \in [u_i, v_i]$, we use convexity of e^x and upper bound it by:

$$e^x \leq \frac{x - u_i}{v_i - u_i} e^{v_i} + \frac{v_i - x}{v_i - u_i} e^{u_i} = \underbrace{\frac{v_i e^{u_i} - u_i e^{v_i}}{v_i - u_i}}_{=b_i} + x \underbrace{\frac{e^{v_i} - e^{u_i}}{v_i - u_i}}_{=c_i}.$$

On the other hand, we use convexity of $f(x) = e^{\frac{x^2}{2}\alpha_0}$ to lower bound it by $f(x) \geq f(u_i) + f'(u_i)(x - u_i)$. Thus, the left-hand side of (13) is lower bounded by:

$$f(u_i) - f'(u_i)u_i - b_i + x(f'(u_i) - c_i + 1) \geq f(u_i) - f'(u_i)u_i - b_i + \min\{v_i(f'(u_i) - c_i + 1), u_i(f'(u_i) - c_i + 1)\}. \quad (14)$$

In the table below, we present the numerical values of (14) for a set of chosen intervals:

interval $[u_i, v_i]$	lower bound (14)
[1.24, 1.78]	0.017
[0.99, 1.24]	0.003
[0.85, 0.99]	0.001
[0.76, 0.85]	0.0007
[0.7, 0.76]	0.001
[0.65, 0.7]	0.0008
[0.6, 0.65]	0.0002
[0.56, 0.6]	0.0008
[0.52, 0.56]	0.0005
[0.47, 0.52]	0.0002
[0.42, 0.47]	0.0004
[0.37, 0.42]	0.0005
[0.34, 0.37]	0.001

As all lower bounds are positive, this finishes the proof.

Note: If we chose $\alpha_0 = 2$, the proof of the lemma would simplify dramatically, as we would only need to separately bound $x \in (-\infty, -1)$, $x \in [-1, 0]$, and $x \in (0, \infty)$. However, we opted for the smallest α_0 , as smaller α_0 translates to a smaller achievable constant in the regret bound. Our choice $\alpha_0 = \frac{9}{8}$ was obtained by performing numerical tests, which showed that this value is very close to the smallest α_0 , for which the inequality still holds.