



HAL
open science

Online optimization and regret guarantees for non-additive long-term constraints

Rodolphe Jenatton, Jim Huang, Cedric Archambeau

► **To cite this version:**

Rodolphe Jenatton, Jim Huang, Cedric Archambeau. Online optimization and regret guarantees for non-additive long-term constraints. 2016. hal-01273728v1

HAL Id: hal-01273728

<https://hal.science/hal-01273728v1>

Preprint submitted on 12 Feb 2016 (v1), last revised 7 Jun 2016 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Online optimization and regret guarantees for non-additive long-term constraints

Rodolphe Jenatton

Amazon.com, Berlin, Germany

JENATTON@AMAZON.COM

Jim C. Huang

Amazon.com, Seattle, WA 98109 USA

HUANGJIM@AMAZON.COM

Cedric Archambeau

Amazon.com, Berlin, Germany

CEDRICA@AMAZON.COM

Abstract

We consider online optimization in the 1-lookahead setting, where the objective does not decompose additively over the rounds of the online game. The resulting formulation enables us to deal with non-stationary and/or long-term constraints, which arise, for example, in online display advertising problems. We propose an online primal-dual algorithm for which we obtain dynamic cumulative regret guarantees. They depend on the convexity and the smoothness of the non-additive penalty, as well as terms capturing the smoothness with which the residuals of the non-stationary and long-term constraints vary over the rounds. We conduct experiments on synthetic data to illustrate the benefits of the non-additive penalty and show vanishing regret convergence on live traffic data collected by a display advertising platform in production.

1. Introduction

Online optimization can be viewed as a sequential game where in each round $t \in \{1, \dots, T\}$, we are required to play an action, represented by a vector \mathbf{x}_t , which takes values in a set of actions $\{\mathcal{X}_t : \mathcal{X}_t \subseteq \mathbb{R}^d\}_{t=1}^T$. We then observe a reward $f_t(\mathbf{x}_t)$ as a function of the action we chose in round t . The goal is to generate a sequence of actions such that some measure of performance is for instance maximized over the course of T rounds.

The performance metric typically adopted in the online

learning framework is the *cumulative regret* (Cesa-Bianchi and Lugosi, 2006). Moreover, online learning traditionally assumes that the objective function has an *additive* structure that nicely decomposes as a sum of regrets over the rounds of the optimization. This means that there is no coupling of \mathbf{x}_t 's across successive rounds when the sets \mathcal{X}_t are decoupled across time. Hence, the challenge in online learning rather lies in the fact that \mathbf{x}_t must be estimated before having access to the reward function f_t , which corresponds to the so-called *0-lookahead* setting (Buchbinder et al., 2012; Andrew et al., 2013). In the particular case of dynamic regret with additive objective functions (Zinkevich, 2003; Cesa-Bianchi et al., 2012; Hall and Willett, 2013; Jadbabaie et al., 2015), one seeks to analyze how, under various assumptions on f_t and \mathcal{X}_t , the sequence of rewards $f_t(\hat{\mathbf{x}}_t)$ collected in an online fashion compares with the best sequence of rewards $f_t(\mathbf{x}_t^*)$ collected in hindsight. The dynamic regret over T rounds is defined as follows:

$$\frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t^*) - \frac{1}{T} \sum_{t=1}^T f_t(\hat{\mathbf{x}}_t), \quad (1)$$

with $\frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t^*) = \max_{\mathbf{x}_1 \in \mathcal{X}_1, \dots, \mathbf{x}_T \in \mathcal{X}_T} \frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t)$. Unlike static regret analysis (Zinkevich, 2003), where we compare against a *single* best action \mathbf{x}^* in retrospect over all rounds played, the best dynamic comparator, as its name indicates, need not be identical for all $t \in \{1, \dots, T\}$.

In this work, we consider a setting similar to online learning with dynamic regret, but we do not assume that successive actions \mathbf{x}_t are decoupled. We provide novel regret bounds for the case where rewards are additive and the regrets are non-additive as the total cost over T rounds is non-decomposable. There is a small body of recent work that study non-additive regrets. For example, Rakhlin et al. (2010) consider a wide class of non-additively decomposable objective functions in a 0-lookahead setting, cover-

ing for instance the problems of Blackwell’s approachability, the calibration of forecasters and the global cost online learning game from Even-Dar et al. (2009). More recently, Kar et al. (2014) handle some specific form of non-decomposability in relation with the online optimization of metrics such as the precision at k . Closer in spirit to our work is the one by Agrawal and Devanur (2015), who provide expected regret bounds in the case where non-additive costs depend only on the empirical mean of the actions.

Metrical task systems (MTS) (Borodin et al., 1992) provide an alternative analysis framework for online optimization. They consider *movement costs* that penalize variations of \mathbf{x}_t ’s *across time* and rely on competitive analysis (Borodin and El-Yaniv, 2005). Similar to online learning with dynamic regret, the performance of the online optimization algorithm is compared to the best sequence of actions. However, instead of measuring cumulative regret, competitive analysis adopts a *multiplicative* metric, known as the competitive ratio. We refer the interested reader to (Buchbinder et al., 2012; Andrew et al., 2013) for a detailed discussion. Moreover, the online game in MTS differs from the online learning setup in that it follows the *1-lookahead* setting where the player has access to the reward function f_t *before* estimating \mathbf{x}_t . The typical instantiation of the movement cost in MTS is a total variation penalty $\sum_{t=1}^{T-1} \|\mathbf{x}_t - \mathbf{x}_{t+1}\|_2$, that can also be defined with non-Euclidean norms (Bera et al., 2013). While the movement cost introduces dependencies across actions \mathbf{x}_t in successive rounds, we note that, to the best of our knowledge, previous work only considered movement costs with an additive structure. We depart from this approach by considering non-additive penalties.

The aforementioned frameworks have been successfully applied to derive and analyze online optimization algorithms in for example online routing (Awerbuch and Kleinberg, 2008), process migration of servers (Borodin and El-Yaniv, 2005) and portfolio allocation (Helmbold et al., 1996). However, the assumption that regrets or movement costs have an additive structure is restrictive in practice. In particular, we consider the problem of *online ad allocation*, which is at the core of modern display advertising systems. The online ad allocation (Chen et al., 2011) problem consists of sequentially allocating ad impressions (encoded by \mathbf{x}_t) to a large number of competing ad slots across a large number of websites and mobile apps, subject to a variety of advertiser objectives and constraints. Advertisers will typically expect that in solving the ad allocation problem, we maximize a measure of ad performance (or advertiser welfare), subject to constraints on user targeting and constraints on ad delivery (e.g., spend as close to 100% of an ad’s budget as possible over T rounds).

Hence, the online ad allocation problem can be viewed as

a hybrid between online learning with dynamic regret and MTS. It falls into the 1-lookahead setting and requires non-additive constraints, but there is no practically-justifiable concept of movement cost on $\hat{\mathbf{x}}_t$ that can be applied to account for non-additive constraints. Moreover, the movement cost is not a sensible penalty to capture the fact that one would like to show different ads in successive rounds. Finally, online ad allocation is concerned with satisfying *long-term* constraints, which means that the cumulative constraint violations resulting from the sequence of vectors $\{\mathbf{x}_t\}_{t=1}^T$ should not exceed a certain amount by the final round T . The previous approaches proposed in the online learning literature (Mahdavi et al., 2012; Agrawal and Devanur, 2015) are not suitable for our use-case since they handle stationary long-term constraints with static regret guarantees in the 0-lookahead setting. And again, this type of constraints is difficult to encode as movement costs.

Contributions. Within the 1-lookahead setting, we study an online optimization problem with dynamic regret defined in terms of a sum of concave rewards and a penalty that does not decompose additively over the rounds of the online game. Non-additive penalties are suitable for modeling *non-stationary* and/or *long-term* constraints, which are of practical importance in display advertising. The resulting formulation is, to the best of our knowledge, novel and extends the work of Mahdavi et al. (2012) and Agrawal and Devanur (2015). We further provide dynamic cumulative regret guarantees, showing that vanishing regret is driven by the convexity and the smoothness of the non-additive penalty, along with the smoothness with which the residuals vary over the rounds. Hence, the structure of our bound echoes results derived in previous work within the context of 0-lookahead dynamic regret analysis (Zinkevich, 2003). Finally, we illustrate our methodological contribution by conducting experiments on synthetic data to validate the benefits of non-additive penalties and study the convergence of the cumulative regret on live traffic data collected by a display advertising platform in production.

2. Problem statement

Motivated by the practical realities of the online ad allocation problem, we focus on a class of online optimization problems in which the actions played over time must not only maximize some rewards but also minimize *long-term* constraints as penalized by some *non-additive* error term \mathcal{E} . Formally, for any $(\mathbf{x}_1, \dots, \mathbf{x}_T) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_T$, we study the online maximization of objective functions of the form

$$\mathcal{P}(\mathbf{x}_1, \dots, \mathbf{x}_T) \triangleq \frac{1}{T} \sum_{t=1}^T f_t(\mathbf{x}_t) - \mathcal{E} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t \right), \quad (2)$$

where the matrices $\mathbf{A}_t \in \mathcal{A} \subseteq \mathbb{R}^{m \times d}$ and vectors $\mathbf{b}_t \in \mathbb{R}^m$ encode the non-stationary, long-term constraints. Although the sets \mathcal{X}_t 's depend on t , they are assumed to be fully available at each round t and do not take part in the online game.

Our goal is to derive an online algorithm and prove that the sequence $(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_T$ it generates is guaranteed to satisfy a regret bound of the form

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \leq \mathcal{R}_T + \mathcal{S}_e + \mathcal{S}_A, \quad (3)$$

where we have defined the optimal (primal) objective

$$\mathcal{P}^* \triangleq \max_{\mathbf{x}_1 \in \mathcal{X}_1, \dots, \mathbf{x}_T \in \mathcal{X}_T} \mathcal{P}(\mathbf{x}_1, \dots, \mathbf{x}_T) = \mathcal{P}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*). \quad (4)$$

The regret bound in (3) is comprised of three terms, each capturing different aspects of the problem.

The first term \mathcal{R}_T quantifies the contribution due to the application of an online projected subgradient method in the dual problem for (2), as described in Section 3.1.

The second term \mathcal{S}_e quantifies the impact of the smoothness with which the sequences of m -dimensional *error vectors* $\{\mathbf{e}_t^*\}_{t=1}^T$ evolve over time, where $\mathbf{e}_t^* \triangleq \mathbf{A}_t \mathbf{x}_t^* - \mathbf{b}_t$. In this respect, the term \mathcal{S}_e will be reminiscent of the guarantees traditionally obtained with dynamic regret analysis (Cesa-Bianchi et al., 2012), with the key exception that the smoothness is not based on the variables \mathbf{x}_t^* themselves, but rather on the error vectors \mathbf{e}_t^* . This term also has a natural interpretation and relevance in the online advertising setting, as discussed in Section 4.2. The derivation for \mathcal{S}_e will be the topic of Section 3.2.

The third and last term \mathcal{S}_A models the impact of having to estimate constraint matrices \mathbf{A}_t via $\hat{\mathbf{A}}_t$, as in practice \mathbf{A}_t is not known at the time when the action \mathbf{x}_t is played. This aspect is also motivated by practical aspects of the online ad allocation problem. As it will be made clear in Section 3.3, \mathcal{S}_A will depend on how smoothly the constraint matrices \mathbf{A}_t 's evolve over time. Given the above, our methodological contributions can thus be viewed as both a regret analysis with dynamic comparators for non-additive objective functions, and a competitive analysis evaluated through a regret criterion – or more precisely, a *competitive difference* criterion, as defined in Andrew et al. (2013) – where the service cost is f_t and our movement cost \mathcal{E} is convex and non-additive. We will interchangeably refer to both competitive difference and dynamic regret in the sequel.

3. Analysis

We describe in this section the main components of our analysis, starting with the saddle point formulation. The proofs of the results are relegated to the appendix.

3.1. Linearization and Lagrangian formulation

Since the main challenge of our analysis lies in the fact that \mathcal{E} breaks the additive structure of (1), a natural strategy is therefore to linearize \mathcal{E} via Fenchel conjugacy, similar to (Agrawal and Devanur, 2015). We define the Fenchel conjugate of \mathcal{E} as $\mathcal{E}^*(\boldsymbol{\lambda}) \triangleq \sup_{\mathbf{z} \in \text{dom}(\mathcal{E})} \{\boldsymbol{\lambda}^\top \mathbf{z} - \mathcal{E}(\mathbf{z})\}$ (Boyd and Vandenberghe, 2004), and from now on, we assume that

(A1) The function \mathcal{E} is convex.¹

(A2) The Fenchel conjugate \mathcal{E}^* of \mathcal{E} has its domain

$$\Lambda \triangleq \text{dom}(\mathcal{E}^*) = \{\boldsymbol{\lambda} \in \mathbb{R}^m : \mathcal{E}^*(\boldsymbol{\lambda}) < +\infty\}$$

that is compact so that there exists $R_\lambda \triangleq \max_{\boldsymbol{\lambda} \in \Lambda} \|\boldsymbol{\lambda}\|_2 < +\infty$.

(A3) For any $t \in \{1, \dots, T\}$, the function f_t is concave, and the set \mathcal{X}_t is compact and convex with $\mathcal{X}_t \subseteq \text{dom}(f_t)$, so there exists $R_x \triangleq \max_{t \in \{1, \dots, T\}, \mathbf{x}_t \in \mathcal{X}_t} \|\mathbf{x}_t\|_2 < +\infty$.

Equipped with those assumptions, we introduce for any $\mathbf{x}_t \in \mathcal{X}_t$, $\boldsymbol{\lambda} \in \Lambda$ the Lagrangian function for round t :

$$\mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) \triangleq f_t(\mathbf{x}_t) - \boldsymbol{\lambda}^\top (\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t) + \mathcal{E}^*(\boldsymbol{\lambda}). \quad (5)$$

It can be observed that $\min_{\boldsymbol{\lambda} \in \Lambda} \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}) = \mathcal{P}(\mathbf{x}_1, \dots, \mathbf{x}_T)$.

In the sequel, we shall refer to $\boldsymbol{\lambda} \in \mathbb{R}^m$ as dual variables. Following previous work on online saddle point optimization (Mahdavi et al., 2012; Koppel et al., 2014), our procedure will alternate between maximizing the Lagrangian with respect to the primal variable \mathbf{x}_t and minimizing with respect to the dual variable $\boldsymbol{\lambda}$. On the primal side, we note that we are interested in the *1-lookahead* setting, where we have access to \mathcal{L}_t before computing our estimate $\hat{\mathbf{x}}_t$, whereas on the dual side, our sequence of $\hat{\boldsymbol{\lambda}}_t$'s will be generated according to the *0-lookahead* setting. This follows the sequencing of events in practice for display advertising, where the dual variables can only be updated after we observe \mathbf{A}_t , which is only observed after we play $\hat{\mathbf{x}}_t$. Before summarizing the procedure in Algorithm 1, we specify a key computational assumption with respect to the primal variables:

(A4) For any $\boldsymbol{\lambda} \in \Lambda$, we can efficiently compute $\arg\max_{\mathbf{x}_t \in \mathcal{X}_t} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda})$.

In this paper, we will be primarily interested in problems where the above assumption (A4) holds (e.g., online ad allocation problems consisting of maximizing linear functions over simplices), so that the complexity and tractability of our proposed algorithm will mostly depend

¹When we refer to convex/concave functions, we implicitly assume closed proper convex/concave functions (Boyd and Vandenberghe, 2004).

Algorithm 1 Online saddle point optimization

Input: Initial dual variable $\hat{\lambda}_1$, step sizes η_t
for $t = 1$ **to** T **do**
 Receive \mathcal{L}_t // 1-lookahead setting
 Compute $\hat{x}_t \in \operatorname{argmax}_{x_t \in \mathcal{X}_t} \mathcal{L}_t(x_t, \hat{\lambda}_t)$
 Update $\hat{\lambda}_{t+1} = \Pi_\Lambda[\hat{\lambda}_t - \eta_t \nabla_\lambda \mathcal{L}_t(\hat{x}_t, \hat{\lambda}_t)]$
end for

on structural properties of f_t and \mathcal{X}_t . Having laid out assumptions, we present our online algorithm in Algorithm 1, which makes use of an online projected subgradient method (OPSM) (Zinkevich, 2003; Hazan et al., 2007) with respect to the dual variables. Interestingly, while the analysis of the primal objective (2) requires a comparison to T different optimal primal variables (i.e., measuring performance via dynamic regret), the online optimization with respect to the dual variables reduces to a static regret analysis that compares with an unique optimal dual variable λ^* . A related observation was exploited by (Shalev-Shwartz and Singer, 2006; Shalev-Shwartz and Kakade, 2009), but following an opposite route, i.e., analyzing the primal objective using static regret, while leveraging the dynamic regret structure of the dual optimization problem. Given Algorithm 1, we now study the impact of applying OPSM for our problem in Lemma 1, subject to an additional assumption about the boundedness of the subgradients of $\mathcal{L}_t(x_t, \cdot)$:

(A5) There exists $G > 0$, $\max_{\substack{t \in \{1, \dots, T\}, \\ x_t \in \mathcal{X}_t, \lambda \in \Lambda}} \|\nabla_\lambda \mathcal{L}_t(x_t, \lambda)\|_2 \leq G$.

Lemma 1. *Let assumptions (A1)-(A3) and (A5) hold. Let $\sigma \geq 0$ be the strong convexity parameter of $\mathcal{L}_t(\hat{x}_t, \cdot)$ with respect to the ℓ_2 norm. It holds for any $\lambda \in \Lambda$, and sequences $\{\hat{x}_t, \hat{\lambda}_t\}_{t=1}^T$ generated by Algorithm 1 that*

$$\frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\hat{x}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{x}_t, \lambda) \leq \mathcal{R}_T,$$

where if $\mathcal{L}_t(\hat{x}_t, \cdot)$ is strongly convex ($\sigma > 0$), we take the step size $\eta_t = \frac{1}{\sigma t}$ and we have $\mathcal{R}_T = \frac{G^2}{2\sigma T} \log(eT)$, while if $\mathcal{L}_t(\hat{x}_t, \cdot)$ is convex ($\sigma = 0$), we can choose the step size $\eta_t = \frac{2R_\lambda}{G\sqrt{T}}$ leading to $\mathcal{R}_T = \frac{2R_\lambda G}{\sqrt{T}}$.

The two claims can be found respectively in (Zinkevich, 2003; Hazan et al., 2007), a short proof is given in the appendix for self-containedness. For the convex case, the step size η_t depends on the time horizon T , and we can use the doubling trick (e.g., see Section 2.3.1 in (Shalev-Shwartz, 2011)) to remove this dependency.

Having presented the component of the regret bound that comes from the use of OPSM, we next examine why the contribution of \mathcal{R}_T is not sufficient by itself to control the regret bound (3).

3.2. Controlling the worst sequence of dual variables

Given Assumptions (A1)-(A3), there do exist some optimal primal variables $\{x_t^*\}_{t=1}^T$ as defined in (4). Let us denote an optimal dual variable given the primal variables $\{x_t^*\}_{t=1}^T$ as

$$\lambda^* \in \operatorname{argmin}_{\lambda \in \Lambda} \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(x_t^*, \lambda),$$

and define $\hat{\lambda}$ equivalently for the sequence $\{\hat{x}_t\}_{t=1}^T$ where the existence of both λ^* and $\hat{\lambda}$ are guaranteed by Assumptions (A1)-(A2).

From the perspective of deriving an upper bound on $\mathcal{P}^* - \mathcal{P}(\hat{x}_1, \dots, \hat{x}_T)$, we observe that $\mathcal{L}_t(x_t^*, \lambda^*) - \mathcal{L}_t(\hat{x}_t, \hat{\lambda})$ can be rewritten as

$$(\mathcal{L}_t(x_t^*, \lambda^*) - \mathcal{L}_t(x_t^*, \hat{\lambda}_t)) + (\mathcal{L}_t(x_t^*, \hat{\lambda}_t) - \mathcal{L}_t(\hat{x}_t, \hat{\lambda})),$$

which in turn, using (A4), is upper-bounded by

$$\underbrace{(\mathcal{L}_t(x_t^*, \lambda^*) - \mathcal{L}_t(x_t^*, \hat{\lambda}_t))}_{\triangleq \mathcal{G}_t(\hat{\lambda}_t)} + \underbrace{(\mathcal{L}_t(\hat{x}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{x}_t, \hat{\lambda}))}_C, \quad (6)$$

where the second term above is readily obtained by the definition of $\hat{x}_t \in \operatorname{argmax}_{x_t \in \mathcal{X}_t} \mathcal{L}_t(x_t, \hat{\lambda}_t)$. This observation is useful in several respects: while the term C can be directly bounded as a result of Lemma 1, the term $\mathcal{G}_t(\hat{\lambda}_t)$, which refers to the gap we incur as a result of the sequentially-generated $(\hat{\lambda}_1, \dots, \hat{\lambda}_T)$, is not accounted for by Lemma 1, and so it needs to be controlled differently. Notably, $\hat{\lambda}_t \mapsto \mathcal{G}_t(\hat{\lambda}_t)$ is concave over Λ , and so is $(\hat{\lambda}_1, \dots, \hat{\lambda}_T) \mapsto \frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\hat{\lambda}_t)$ over Λ^T : this implies that we can cast the problem of controlling the worst sequence of dual variables $(\hat{\lambda}_1, \dots, \hat{\lambda}_T)$ as a concave maximization problem.

In order to obtain a meaningful upper bound on $\mathcal{G}_t(\hat{\lambda}_t)$ and $\frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\hat{\lambda}_t)$, we leverage the fact that we are not dealing with any general sequence $(\hat{\lambda}_1, \dots, \hat{\lambda}_T) \in \Lambda^T$, but sequences possibly output by Algorithm 1. More precisely, we make this characterization by noting that OPSM generates successive $\hat{\lambda}_t$ and $\hat{\lambda}_{t+1}$ estimates whose differences $\|\hat{\lambda}_t - \hat{\lambda}_{t+1}\|_2$ are controlled as a function of the step size η_t (Andrew et al., 2013). Following (Zinkevich, 2003), we therefore introduce the convex set

$$\Lambda_{T,\varepsilon} \triangleq \left\{ (\lambda_1, \dots, \lambda_T) \in \Lambda^T : \sum_{t=1}^{T-1} \|\lambda_t - \lambda_{t+1}\|_2 \leq \varepsilon \right\}, \quad (7)$$

which will be useful shortly in deriving an upper bound for \mathcal{S}_e . In particular, assuming (A5) holds and setting $\varepsilon = G \sum_{t=1}^T \eta_t$, for any sequence $(\hat{\lambda}_1, \dots, \hat{\lambda}_T) \in \Lambda^T$ generated by Algorithm 1, we have $(\hat{\lambda}_1, \dots, \hat{\lambda}_T) \in \Lambda_{T,\varepsilon}$ and

$$\frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\hat{\lambda}_t) \leq \mathcal{S}_e \triangleq \max_{(\lambda_1, \dots, \lambda_T) \in \Lambda_{T,\varepsilon}} \frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\lambda_t). \quad (8)$$

We next turn to some lemmas that make the expression of \mathcal{S}_e more explicit by leveraging duality arguments. To this end, we introduce some additional notation: first, we rewrite the total-variation constraint from (7) as

$$\sum_{t=1}^{T-1} \|\lambda_t - \lambda_{t+1}\|_2 = \Omega_{1/2}((\Delta \otimes \mathbf{I})\underline{\lambda})$$

where $\underline{\lambda} \in \mathbb{R}^{T \cdot m}$ stands for the vector formed by stacking the m -dimensional λ_t 's, \otimes is the Kronecker product, $\Delta \in \mathbb{R}^{(T-1) \times T}$ is the discrete 1-dimensional gradient matrix,² while $\Omega_{1/2}$ is the ℓ_1/ℓ_2 norm with dual norm $\Omega_{\infty/2}$, such that

$$\Omega_{1/2}(\mathbf{u}) \triangleq \sum_{t=1}^{T-1} \|\mathbf{u}_t\|_2 \quad \text{and} \quad \Omega_{\infty/2}(\mathbf{v}) \triangleq \max_{t \in \{1, \dots, T-1\}} \|\mathbf{v}_t\|_2,$$

for any vectors $\mathbf{u} = [\mathbf{u}_1, \dots, \mathbf{u}_{T-1}]$ and $\mathbf{v} = [\mathbf{v}_1, \dots, \mathbf{v}_{T-1}]$ in $\mathbb{R}^{(T-1) \cdot m}$. We now show that \mathcal{S}_e can conveniently be expressed as a minimization problem.

Lemma 2. *Let $\mathbf{e}_t^* = \mathbf{A}_t \mathbf{x}_t^* - \mathbf{b}_t$. It holds that*

$$\mathcal{S}_e = \min_{\alpha \in \mathbb{R}^{(T-1) \cdot m}} \mathcal{S}_e(\alpha)$$

where we have introduced the convex function $\alpha \mapsto \mathcal{S}_e(\alpha)$

$$\frac{1}{T} \sum_{t=1}^T \mathcal{E}(\mathbf{e}_t^* - [(\Delta \otimes \mathbf{I})^\top \alpha]_t) - \mathcal{E}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{e}_t^*\right) + \frac{\varepsilon}{T} \Omega_{\infty/2}(\alpha).$$

The proof is given in the appendix and relies on duality arguments. We note at this juncture that our analysis can similarly handle other penalties related to the total-variation chosen in (7), for instance $\|(\Delta \otimes \mathbf{I})\underline{\lambda}\|_F^2$ (we omit the details owing to space limitations).

The expression for \mathcal{S}_e provided by Lemma 2 (the result of a *minimization* problem), makes it possible to obtain an upper bound for any candidate vector $\alpha \in \mathbb{R}^{(T-1) \cdot m}$. We propose below one such instantiation that highlights how \mathcal{S}_e depends on the smoothness with which the sequence $\{\mathbf{e}_t^*\}_{t=1}^T$ varies over time, subject to the following assumptions on \mathcal{E} :

(A7) \mathcal{E} is sub-additive: there exists $c \geq 1$ such that for any two vectors $\mathbf{u}, \mathbf{v} \in \text{dom}(\mathcal{E})$, $\mathcal{E}(\mathbf{u} + \mathbf{v}) \leq c[\mathcal{E}(\mathbf{u}) + \mathcal{E}(\mathbf{v})]$,

(A8) \mathcal{E} is symmetric around 0: For any $\mathbf{u} \in \text{dom}(\mathcal{E})$, it holds that $-\mathbf{u} \in \text{dom}(\mathcal{E})$ and $\mathcal{E}(\mathbf{u}) = \mathcal{E}(-\mathbf{u})$.

Lemma 3. *Let Assumptions (A7), (A8) hold and let $\mathbf{e}_t^* = \mathbf{A}_t \mathbf{x}_t^* - \mathbf{b}_t$. The term \mathcal{S}_e is upper-bounded by*

$$\frac{2c}{T} \sum_{t=1}^{T-1} \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*) - \mathcal{E}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{e}_t^*\right) + \frac{\varepsilon}{T} \Omega_{\infty/2}((\Delta \otimes \mathbf{I})\underline{\mathbf{e}}^*),$$

where $\underline{\mathbf{e}}^* \in \mathbb{R}^{T \cdot m}$ stands for the vector with stacked \mathbf{e}_t^* 's.

²The matrix $\Delta \in \mathbb{R}^{(T-1) \times T}$ contains two non-zero entries per row, with $\Delta_{t,t} = 1$ and $\Delta_{t,t+1} = -1$ for $t \in \{1, \dots, T-1\}$.

The proof of this lemma can be found in the appendix. It is worth mentioning that Assumptions (A7)-(A8) are not strictly necessary to obtain upper-bounds on \mathcal{S}_e , but are invoked in the above lemma in order to simplify the resulting expression. We can see from Lemma 3 that \mathcal{S}_e captures both the cumulative variations of $\{\mathbf{e}_t^*\}_{t=1}^T$ through $\sum_{t=1}^{T-1} \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*)$ and the worst of those variations via $\Omega_{\infty/2}((\Delta \otimes \mathbf{I})\underline{\mathbf{e}}^*) = \max_{t \in \{1, \dots, T-1\}} \|\mathbf{e}_t^* - \mathbf{e}_{t+1}^*\|_2$ modulated by the average of the step sizes $\varepsilon/T = G \sum_{t=1}^T \eta_t/T$. The structure of this upper-bound is analogous to that developed in (Zinkevich, 2003) in the context of additive regret objectives and dynamic comparators.

Having derived terms \mathcal{R}_T and \mathcal{S}_e in the cumulative regret bound, we now turn to the description of the last term \mathcal{S}_A of our regret guarantee (3).

3.3. Estimating the matrices $\{\mathbf{A}_t\}_{t=1}^T$

So far we have assumed that we have access to \mathbf{A}_t at each round, when in practice we only get to observe \mathbf{A}_t after we take an action $\hat{\mathbf{x}}_t$. For example, in the online ad allocation problem, the amount of money to be charged to the advertiser for a single ad impression is revealed only after we have made a decision on which ad to allocate for the impression, i.e., once we have computed $\hat{\mathbf{x}}_t$. We now address the cost incurred in having to estimate the constraint matrices \mathbf{A}_t before playing $\hat{\mathbf{x}}_t$. To this end, we assume that in addition to the dual variables being bounded with radius R_λ , we have

(A9) The set \mathcal{A} is convex and bounded; in particular, there exists $R_A \triangleq \max_{\mathbf{A} \in \mathcal{A}} \|\mathbf{A}\|_F < +\infty$.

Since the choice of $\hat{\mathbf{x}}_t$ now depends on the estimate $\hat{\mathbf{A}}_t$, we introduce the estimated Lagrangian

$$\hat{\mathcal{L}}_t(\mathbf{x}_t, \lambda) \triangleq f_t(\mathbf{x}_t) - \lambda^\top (\hat{\mathbf{A}}_t \mathbf{x}_t - \mathbf{b}_t) + \mathcal{E}^*(\lambda), \quad (9)$$

with the direct relationship (for any $\mathbf{x}_t \in \mathcal{X}_t, \lambda \in \Lambda$):

$$\hat{\mathcal{L}}_t(\mathbf{x}_t, \lambda) = \mathcal{L}_t(\mathbf{x}_t, \lambda) - \lambda^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t) \mathbf{x}_t. \quad (10)$$

Combining the above expression (10) along with the decomposition detailed in (6), it can be shown (see Lemma A in the appendix) that the residual term due to the estimation of \mathbf{A}_t is given by

$$\mathcal{S}_A \triangleq \frac{1}{T} \sum_{t=1}^T \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t) (\mathbf{x}_t^* - \hat{\mathbf{x}}_t). \quad (11)$$

As a result of having to estimate \mathbf{A}_t , we present Algorithm 2, where we also apply an OPSM to estimate the constraint matrices \mathbf{A}_t 's.

Lemma 4 then quantifies the cumulative regret incurred as a result of having to use estimated constraint matrices $\hat{\mathbf{A}}_t$

Algorithm 2 Online saddle point optimization with estimated constraint matrices $\hat{\mathbf{A}}_t$

Input: Initial $\hat{\mathbf{A}}_1$ and dual variable $\hat{\lambda}_1$, step sizes η_t, ν_t
for $t = 1$ **to** T **do**
 Receive $\hat{\mathcal{L}}_t$ // 1-lookahead and $\hat{\mathbf{A}}_t$
 Compute $\hat{\mathbf{x}}_t \in \operatorname{argmax}_{\mathbf{x}_t \in \mathcal{X}_t} \hat{\mathcal{L}}_t(\mathbf{x}_t, \hat{\lambda}_t)$
 Receive \mathbf{A}_t
 Update $\hat{\lambda}_{t+1} = \Pi_\Lambda[\hat{\lambda}_t - \eta_t \nabla_{\lambda} \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}_t)]$
 // \mathbf{G}_t a subgradient of $\|\mathbf{A}_t - \mathbf{A}\|_{\mathbb{F}}$
 Update $\hat{\mathbf{A}}_{t+1} = \Pi_{\mathcal{A}}[\hat{\mathbf{A}}_t - \nu_t \mathbf{G}_t]$
end for

to produce $\hat{\mathbf{x}}_t$, where the result stems from the analysis of (Hall and Willett, 2013).

Lemma 4. Consider Algorithm 2. For OPSM on matrices $\hat{\mathbf{A}}_t$ with step size $\nu_t = R_A/\sqrt{t}$, we have

$$\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{A}}_t - \mathbf{A}_t\|_{\mathbb{F}} \leq \frac{3}{\sqrt{T}} \left[R_A + \sum_{t=1}^T \|\mathbf{A}_t - \mathbf{A}_{t+1}\|_{\mathbb{F}} \right],$$

and the term $\mathcal{S}_{\mathbf{A}}$ from the general regret bound (3), as defined in (11), is upper-bounded by:

$$\mathcal{S}_{\mathbf{A}} \leq \frac{6R_\lambda R_x}{\sqrt{T}} \left[R_A + \sum_{t=1}^T \|\mathbf{A}_t - \mathbf{A}_{t+1}\|_{\mathbb{F}} \right].$$

Having introduced the main three components of our regret bound, we next formally present our main results.

3.4. Main results and discussion

We start by stating our core theorem:

Theorem 1. Let Assumptions (A1)-(A9) hold, and let $\mathbf{x}_1^*, \dots, \mathbf{x}_T^*$ be a sequence of optimal primal solutions for (2). Define the optimal error vectors $\mathbf{e}_t^* \triangleq \mathbf{A}_t \mathbf{x}_t^* - \mathbf{b}_t$ and $M_e \triangleq \max_{t \in \{1, \dots, T-1\}} \|\mathbf{e}_t^* - \mathbf{e}_{t+1}^*\|_2$. Algorithm 2 generates a sequence $(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \in \prod_{t=1}^T \mathcal{X}_t$ satisfying

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \leq \mathcal{R}_T + \mathcal{S}_e + \mathcal{S}_{\mathbf{A}}$$

where $\mathcal{S}_{\mathbf{A}} = \frac{6R_\lambda R_x}{\sqrt{T}} [R_A + \sum_{t=1}^T \|\mathbf{A}_t - \mathbf{A}_{t+1}\|_{\mathbb{F}}]$ along with $\mathcal{R}_T = \frac{2R_\lambda G}{\sqrt{T}}$ and

$$\mathcal{S}_e = \frac{2c}{T} \sum_{t=1}^{T-1} \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*) - \mathcal{E}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{e}_t^*\right) + M_e \frac{2R_\lambda}{\sqrt{T}}.$$

If we additionally assume that

(A10) \mathcal{E} has Lipschitz continuous gradients over its domain with parameter $L > 0$,

we can take instead $\mathcal{R}_T = \frac{LG^2}{2T} \log(eT)$ and the term

$$M_e \frac{2R_\lambda}{\sqrt{T}} \text{ can be replaced by } M_e \frac{LG \log(eT)}{T}.$$

The proof of the results can be found in the appendix. The statement of Theorem 1 calls for some comments. Omitting the contribution of terms depending on $\{R_\lambda, R_x, R_A, G, M_e\}$, Theorem 1 guarantees that when \mathbf{A}_t is assumed to be known, the regret $\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T)$ is upper-bounded by $\mathcal{O}(\frac{1}{T} \sum_{t=1}^{T-1} \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*))$ and terms scaling with $\mathcal{O}(1/\sqrt{T})$, or $\mathcal{O}(\log(T)/T)$ (for convex or strongly convex cases respectively). As a result, the upper bound is mainly driven by how smooth the sequence $\{\mathbf{e}_t^*\}_{t=1}^T$ varies over time, as measured by \mathcal{E} . In the more challenging setting where the constraint matrices \mathbf{A}_t 's are also estimated, we pay an additional cost $\mathcal{O}(\frac{1}{\sqrt{T}} [1 + \sum_{t=1}^T \|\mathbf{A}_t - \mathbf{A}_{t+1}\|_{\mathbb{F}}])$, so that the potential strong convexity of \mathcal{E}^* (or equivalently, the gradient Lipschitz continuity of \mathcal{E}) plays only a second-order role in the above theorem.

Possible instantiations of \mathcal{E} : We now present possible valid instantiations of the error function \mathcal{E} . A first example for which the set of assumptions (A1)-(A2) can be satisfied considers $\mathcal{E}(\mathbf{z}) = r \cdot \Omega(\mathbf{z})$ where Ω refers to any norm on \mathbb{R}^m . Indeed, it is well known (e.g., see Example 3.26 in (Boyd and Vandenberghe, 2004)) that \mathcal{E}^* corresponds in this case to the indicator function of the ball for the dual norm Ω^* with radius r . In particular, if Ω is the ℓ_2 norm, we can take $R_\lambda = r$.

Tighter regret bounds can be obtained if the function \mathcal{E} is taken to be gradient Lipschitz continuous (see Theorem 1 and Assumption (A10)), while preserving the boundedness of the domain Λ of its Fenchel conjugate. A possible choice in this case is an instance of a Huber function (e.g., see Section 10.6 in (Hastie et al., 2009)), as detailed in Table 1.

In the setting of the online ad allocation problem, it is sometimes required to impose an *asymmetric* error function (e.g., to penalize under- and over-delivery differently). We show in Table 1 corresponding valid instantiations of \mathcal{E} that depend only on their arguments via their positive parts.³

Complexity of computing $\hat{\mathbf{x}}_t$: Assumption (A4), about the exact computational oracle for $\hat{\mathbf{x}}_t$, hides, and concentrates, the difficulties related to the optimization with respect to the primal variables \mathbf{x}_t . We can observe, in the light of the decomposition (6), that an approximate maximization—for instance leveraging the concavity of f_t —would lead to residual terms that would be hard to control in our online setting. Although it may appear at first sight that an exact maximization is an overly strong requirement, wide classes of problems fall within the scope of this assumption; we can for instance cite the exactly-solved subproblems that are commonly encountered in the context of

³Note that the symmetry assumption (A8) does not hold anymore, but it can be shown that we can still obtain an upper bound of \mathcal{S}_e along the lines of Lemma 3, though with a less compact expression.

$\mathcal{E}(\mathbf{z})$	$\mathcal{E}^*(\boldsymbol{\lambda})$	Λ	c
$\mathcal{H}_{R_\lambda, L}(\ \mathbf{z}\ _2)$	$\mathcal{I}_{\mathcal{B}_{R_\lambda}}(\boldsymbol{\lambda}) + \frac{1}{2L}\ \boldsymbol{\lambda}\ _2^2$	\mathcal{B}_{R_λ}	2
$\mathcal{H}_{R_\lambda, L}(\ \mathbf{z}\ _+)$	$\mathcal{I}_{\mathcal{B}_{R_\lambda} \cap \mathbb{R}_+^m}(\boldsymbol{\lambda}) + \frac{1}{2L}\ \boldsymbol{\lambda}\ _2^2$	$\mathcal{B}_{R_\lambda} \cap \mathbb{R}_+^m$	2
$R_\lambda \cdot \ \mathbf{z}\ _2$	$\mathcal{I}_{\mathcal{B}_{R_\lambda}}(\boldsymbol{\lambda})$	\mathcal{B}_{R_λ}	1
$R_\lambda \cdot \ \mathbf{z}\ _+$	$\mathcal{I}_{\mathcal{B}_{R_\lambda} \cap \mathbb{R}_+^m}(\boldsymbol{\lambda})$	$\mathcal{B}_{R_\lambda} \cap \mathbb{R}_+^m$	1

Table 1. Examples of non-additive penalty functions. For simplicity, we display only expressions based on the ℓ_2 norm, but similar formula can be obtained for $\ell_q, q \in \{1, \infty\}$. We recall that $\Lambda \triangleq \text{dom}(\mathcal{E}^*) = \{\boldsymbol{\lambda} : \mathcal{E}^*(\boldsymbol{\lambda}) < +\infty\}$ and $c \geq 1$ stands for the constant of sub-additivity for \mathcal{E} (see Assumption (A7)). We define \mathcal{B}_r as the Euclidean ball in \mathbb{R}^m with radius r , while \mathcal{I}_C is the indicator function of the convex set C . The function $\mathcal{H}_{r,s}$ is defined as $\mathcal{H}_{r,s}(t) \triangleq \frac{1}{2} \min\{st^2, \frac{t^2}{s}\} + r[|t| - \frac{r}{s}]_+$. The expressions are derived in Lemmas B-C in the appendix.

proximal methods (Parikh and Boyd, 2013) or conditional-gradient algorithms (Jaggi, 2013, see Table 1). We conjecture that, thanks to Assumption (A4), we may weaken Assumption (A3) by, for instance, trying to drop the concavity assumption of the f_t 's and the convexity of the \mathcal{X}_t 's.

Lower bounds: Based on Assumptions (A1)-(A2)-(A3), we know that the minimax equality for our saddle point problem holds (Sion et al., 1957). In particular, defining the (convex) dual function

$$\mathcal{D}(\boldsymbol{\lambda}) \triangleq \frac{1}{T} \sum_{t=1}^T \mathcal{D}_t(\boldsymbol{\lambda}) \triangleq \frac{1}{T} \sum_{t=1}^T \max_{\mathbf{x}_t \in \mathcal{X}_t} \mathcal{L}_t(\mathbf{x}_t, \boldsymbol{\lambda}),$$

we have the equality

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) = \min_{\boldsymbol{\lambda} \in \Lambda} \mathcal{D}(\boldsymbol{\lambda}) - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T).$$

This new relationship makes it possible to derive complementary guarantees for our problem. For instance, when we need not estimate the constraint matrices \mathbf{A}_t 's,⁴ it can be shown that we can *lower bound* $\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T)$ by

$$\left[\frac{1}{T} \sum_{t=1}^T \mathcal{D}_t(\hat{\boldsymbol{\lambda}}_t) - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) - \mathcal{R}_T \right]_+,$$

which sets a lower limit on the best performance we could get with the sequences $\{\hat{\mathbf{x}}_t, \hat{\boldsymbol{\lambda}}_t\}_{t=1}^T$. Interestingly, this lower bound can not only be practically computed, but it also gives an indication about the inherent difficulty of the problem at hand (e.g., if it is large, the online strategy cannot compete efficiently with its offline counterpart).

4. Experiments

In this section, we conduct two sets of experiments. First, we consider a synthetic data and demonstrate the benefit

⁴When \mathbf{A}_t 's are also estimated, the lower bound is more involved, and requires the introduction of an estimated dual function $\hat{\mathcal{D}}_t$ along the lines of (9).

of handling truly non-additive penalties over a heuristic based on an additive relaxation. Second, we report results on an online advertising data set consisting of a sample of 3.3 million bid requests gathered from a large ad serving system used in production at Amazon. Hence, these data account for external advertiser constraints encountered in practice, such as user behavioral targeting. We show that the rate at which regret vanishes matches our theory.

4.1. Additive versus non-additive modeling

We first investigate if the use of a non-additive penalty on average constraint violations yields improvements in reward and/or constraint violations compared to the additive error formulation adopted in standard online learning. The additive long-term penalty is defined as follows:

$$\frac{1}{T} \sum_{t=1}^T [f_t(\mathbf{x}_t) - \mathcal{E}(\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t)]. \quad (12)$$

In the experiments, we consider a linear reward function $f_t(\mathbf{x}) = \mathbf{u}_t^\top \mathbf{x}_t$, with the simplicial constraint $\mathbf{x}_t \in \mathcal{X}_t \triangleq \{\mathbf{x}_t \in \mathbb{R}_+^d : \mathbf{1}^\top \mathbf{x}_t \leq 1\}$. These choices of f_t and \mathcal{X}_t are reminiscent of the online ad allocation setting we explore at greater length in Section 4.2. We simulated data with $m = 25, d = 10, T = 200$ and where $\{\mathbf{A}_t, \mathbf{b}_t, \mathbf{u}_t\}$ were generated to be standard random Gaussian matrices and vectors normalized to unit norms. We ran our online algorithm with primal-dual updates dictated by the non-additive problem formulation in (2) and compared it to a baseline algorithm with primal-only updates following a standard online learning formulation with additive penalties on constraint violations as in (12). Since the exact maximization $\max_{\mathbf{x}_t \in \mathcal{X}_t} [f_t(\mathbf{x}_t) - \mathcal{E}(\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t)]$ may not be obtained in closed-form in this case, we used the solver CVXPY (Diamond et al., 2014). We shall refer to the two approaches as respectively Non-additive and Additive.

So as to cover different types of geometries and convexity assumptions, we consider several choices of penalty functions, namely $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_q$ for $q \in \{1, 2, \infty\}$, along with $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \mathcal{H}_{1,1}(\|\mathbf{z}\|_2)$ where $\mathcal{H}_{1,1}$ is the Huber function defined in Table 1. For each \mathcal{E} , and for, both, Non-additive and Additive, we compute the reward over T rounds given by $\frac{1}{T} \sum_t \mathbf{u}_t^\top \hat{\mathbf{x}}_t$ and the (normalized) non-additive penalty $\mathcal{E}(\frac{1}{T} \sum_t \mathbf{A}_t \hat{\mathbf{x}}_t - \mathbf{b}_t) / R_\lambda$ for the primal variables $\hat{\mathbf{x}}_t$ generated online by the two algorithms.

Figure 1 shows the resulting reward versus constraint violation for varying values of R_λ , averaged over 10 generations of $\{\mathbf{A}_t, \mathbf{b}_t, \mathbf{u}_t\}$. As can be seen from the curves, in applications where constraint violation are to be measured as a non-additive penalty, the additive relaxation (12) leads to significant deterioration of the performance compared to the non-additive penalty. In particular, the domain




Figure 1. Reward $\frac{1}{T} \sum_t \mathbf{u}_t^\top \hat{\mathbf{x}}_t$ as a function of non-additive penalty $\mathcal{E}(\frac{1}{T} \sum_t \mathbf{A}_t \hat{\mathbf{x}}_t - \mathbf{b}_t)/R_\lambda$ for $R_\lambda = 2^\gamma$ with $\gamma \in \{-8, -7.5, \dots, 10\}$. Red crosses correspond to our proposed online algorithm, while blue circles stand for a baseline algorithm with additive penalties. Each subplot displays a different instantiation of \mathcal{E} , namely $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_q$ for $q \in \{1, 2, \infty\}$ and $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \mathcal{H}_{1,1}(\|\mathbf{z}\|_2)$ where $\mathcal{H}_{1,1}$ is defined in Table 1 (best seen in color).

of the achievable constraint violation is very narrow when applying the additive heuristic, meaning that there is little room for making trading-offs. As expected, we also observe that points of Non-additive and Additive superimpose in the regime where $R_\lambda \ll 1$, i.e., when the two formulations focus on the optimization of the reward. Finally, we remark that since the quality of the additive relaxation (12) essentially hinges on the gap in the Jensen’s inequality $\mathcal{E}(\frac{1}{T} \sum_{t=1}^T \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t) \leq \frac{1}{T} \sum_{t=1}^T \mathcal{E}(\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t)$, and therefore on the distribution of the residuals $\{\mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t\}_{t=1}^T$, we provide in the appendix additional simulations where $\{\mathbf{A}_t, \mathbf{b}_t, \mathbf{u}_t\}$ are generated according to different distributions, viz, Cauchy, uniform and gamma. In a nutshell, the same conclusions hold, even for the distributions that appear to make the relaxation (12) tighter.

4.2. Regret convergence on real online advertising data

In the second experiment, we focus on the online ad allocation problem in display advertising subject to long-term constraints on budget consumed per ad. More formally, in the display advertising setting, each user visit to a website or app triggers a bid request i , where each bid request has some subset of the m possible ads that can be served. Indexing ads by $j \in \{1, \dots, m\}$, the welfare of serving an impression for ad j (i.e., the value of showing the ad once to

a user) for bid request i is given by $u_{ij} \geq 0$. For bid request i , let primal variable $x_{ij} = 1$ correspond to our decision to allocate ad j an impression, with $x_{ij} = 0$ otherwise.

To map this problem to our formulation, we will further partition the time dimension into disjoint intervals $t \in \{1, \dots, T\}$ where the number of bid requests in each interval is equal to N . We let I_t be the set of bid requests in round t with $|I_t| = N$. Thus, for a given t , the primal variable \mathbf{x}_t will denote the flattened matrix x_{ij} for $i \in I_t$ and for all ads $j \in \{1, \dots, m\}$, so that $\mathbf{x}_t \in [0, 1]^d$ with $d = N \cdot m$. Since we can show at most one ad per bid request, we additionally have the simplicial constraints $\sum_{j=1}^m x_{ij} \leq 1$ for each bid request i , and we therefore define the set $\mathcal{X}_t \subseteq [0, 1]^d$ as the Cartesian product of those simplices for each $i \in I_t$. We define the welfare vector $\mathbf{u}_t \in \mathbb{R}^d$ following the same flattening operation, so that $f_t(\mathbf{x}_t) = \mathbf{u}_t^\top \mathbf{x}_t$. Matrix $\mathbf{A}_t \in \mathbb{R}^{m \times d}$ is such that the j -th entry $[\mathbf{A}_t \mathbf{x}_t]_j$ represents the amount of budget consumed for ad j as a result of the allocation vector \mathbf{x}_t under a cost per impression model. Finally, the non-additive long-term constraints are modeled as $\frac{1}{T} \sum_t \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t$. Vector \mathbf{b}_t is equal to the constant vector \mathbf{b} of ad budgets. Hence, the penalty on $\frac{1}{T} \sum_t \mathbf{A}_t \mathbf{x}_t - \mathbf{b}_t$ penalizes deviations from spending 100% of each ad’s budget.

The goal of the experiment is to compare the regret behavior of our proposed online algorithm for different choices of non-additive \mathcal{E} on the online advertising data set from live traffic. For each bid request in our data set, each eligible ad candidate comes with a pre-defined welfare and cost to the advertiser, which determine \mathbf{u}_t and \mathbf{A}_t respectively. We partitioned the 3.3 million impressions into batches of equal size, for a total of $T = 10000$ rounds. We computed the average cumulative regret $\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T)$ as a function of T for (symmetric) \mathcal{E} chosen to be i) $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_1$ referred to as `Convex` and ii) $\mathcal{E}(\mathbf{z}) = \mathcal{H}_{R_\lambda, 1}(\|\mathbf{z}\|_2)$ referred to as `Strongly convex`. For all problem instantiations, we set $R_\lambda = 50000$. We computed the term $\mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T)$ by running our online primal-dual algorithm over 20 permutations of bid requests. The term \mathcal{P}^* was obtained using an offline primal-dual method where the dual variable updates correspond to an offline subgradient (resp. gradient) method for the convex (resp. strongly convex) \mathcal{E}^* . Figure 2 shows the resulting cumulative regret as a function of T , averaged over the 20 runs (with error bars that are negligible). We observe that the cumulative regrets incurred for, both, `Convex` and `Strongly convex` decrease as T increases, with the latter decreasing at faster rate than the former as predicted by the theory.

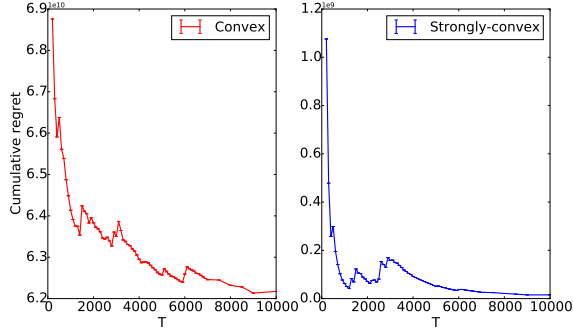


Figure 2. Cumulative regret, averaged over 20 runs, measured on the data collected by a real ad serving system as a function of the number of rounds T for $\mathcal{E}(\mathbf{z}) = R_\lambda \|\mathbf{z}\|_1$ (red) and $\mathcal{E}(\mathbf{z}) = \mathcal{H}_{R_\lambda, 1}(\|\mathbf{z}\|_2)$ (blue). Figure best seen in color.

5. Appendix: Proofs

5.1. Proofs of Lemma 1

Proof. The two claims of the lemma can be found respectively in (Zinkevich, 2003; Hazan et al., 2007) and follows from (possibly, strong) convexity of $\mathcal{L}_t(\hat{\mathbf{x}}_t, \boldsymbol{\lambda})$ with respect to $\boldsymbol{\lambda}$. For completeness, we briefly repeat the core arguments from (Zinkevich, 2003; Hazan et al., 2007). The conclusions come from the fact that the term

$$\sum_{t=1}^T \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\boldsymbol{\lambda}}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \boldsymbol{\lambda})$$

is shown to be upper bounded by, with $a_t \triangleq \|\boldsymbol{\lambda} - \hat{\boldsymbol{\lambda}}_t\|_2^2$,

$$\frac{1}{2} \sum_{t=1}^T \left[\frac{1}{\eta_t} (a_t - a_{t+1}) - \sigma a_t + \eta_t G^2 \right]$$

in turn upper bounded by

$$\frac{a_1}{2} \left[\frac{1}{\eta_1} - \sigma \right] + \frac{G^2}{2} \sum_{t=1}^T \eta_t + \frac{1}{2} \sum_{t=2}^T a_t \left[\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} - \sigma \right].$$

The choices of the step sizes lead to the advertised instantiations of \mathcal{R}_T . \square

5.2. Proofs of Lemma 2

Proof. The result stems from an application of strong duality for the above convex program. First, notice that

$$\frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\boldsymbol{\lambda}_t) = \frac{1}{T} \sum_{t=1}^T \boldsymbol{\lambda}_t^\top \mathbf{e}_t^* - \mathcal{E}^*(\boldsymbol{\lambda}_t) - \mathcal{E}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{e}_t^*\right).$$

Introduce the equality constraint $\mathbf{a} = (\boldsymbol{\Delta} \otimes \mathbf{I})\boldsymbol{\lambda}$, and write the corresponding Lagrangian (we momentarily omit the

constant terms independent of $\boldsymbol{\lambda}_t$ that have no effects on the maximization), that is, for any $\boldsymbol{\lambda}_t \in \Lambda$, $\mathbf{a} \in \mathbb{R}^{(T-1) \cdot m}$, $\boldsymbol{\alpha} \in \mathbb{R}^{(T-1) \cdot m}$ and $\beta \geq 0$:

$$\begin{aligned} L(\boldsymbol{\lambda}, \mathbf{a}, \boldsymbol{\alpha}, \beta) &= \frac{1}{T} \sum_{t=1}^T \boldsymbol{\lambda}_t^\top \mathbf{e}_t^* - \mathcal{E}^*(\boldsymbol{\lambda}_t) \\ &\quad + \boldsymbol{\alpha}^\top (\mathbf{a} - (\boldsymbol{\Delta} \otimes \mathbf{I})\boldsymbol{\lambda}) + \beta(\varepsilon - \Omega_{1/2}(\mathbf{a})). \end{aligned}$$

Using the fact that \mathcal{E}^* is the Fenchel conjugate of \mathcal{E} and that the conjugate of the l_1/l_2 norm $\Omega_{1/2}$ is the indicator function of the unit ball for its dual l_∞/l_2 norm (Example 3.26 in (Boyd and Vandenberghe, 2004)), we can maximize out the primal variables:

$$\max_{\boldsymbol{\lambda}, \mathbf{a}} L(\boldsymbol{\lambda}, \mathbf{a}, \boldsymbol{\alpha}, \beta) = \frac{1}{T} \sum_{t=1}^T \mathcal{E}\left(\mathbf{e}_t^* - [T(\boldsymbol{\Delta} \otimes \mathbf{I})^\top \boldsymbol{\alpha}]_t\right) + \beta \varepsilon$$

under the constraint $\max_{t=1 \dots T-1} \|\boldsymbol{\alpha}_t\|_2 \leq \beta$. Minimizing out β , and making the change of variable $\boldsymbol{\alpha} \leftarrow T \cdot \boldsymbol{\alpha}$ leads to the claimed dual minimization problem. Finally, the initial claim holds with equality, since Slater's constraint qualification applies here to guarantee strong duality (Section 5.2.3 in (Boyd and Vandenberghe, 2004)). \square

5.3. Proofs of Lemma 3

Proof. We evaluate the minimization problem $\min_{\boldsymbol{\alpha} \in \mathbb{R}^{(T-1) \cdot m}} \mathcal{S}_e(\boldsymbol{\alpha})$ at $\boldsymbol{\alpha} = -(\boldsymbol{\Delta} \otimes \mathbf{I})\underline{\mathbf{e}}^*$ where $\underline{\mathbf{e}}^* \in \mathbb{R}^{T \cdot d}$ is the vector of stacked \mathbf{e}_t^* 's. First notice that

$$\begin{aligned} \mathbf{e}_t^* - [(\boldsymbol{\Delta} \otimes \mathbf{I})^\top \boldsymbol{\alpha}]_t &= \mathbf{e}_t^* + [(\boldsymbol{\Delta}^\top \boldsymbol{\Delta} \otimes \mathbf{I})\underline{\mathbf{e}}^*]_t \\ &= [((\mathbf{I} + \boldsymbol{\Delta}^\top \boldsymbol{\Delta}) \otimes \mathbf{I})\underline{\mathbf{e}}^*]_t, \end{aligned}$$

and define $\mathbf{w} \triangleq ((\mathbf{I} + \boldsymbol{\Delta}^\top \boldsymbol{\Delta}) \otimes \mathbf{I})\underline{\mathbf{e}}^*$. We then get, by definition of \mathcal{S}_e ,

$$\mathcal{S}_e \leq \frac{1}{T} \sum_{t=1}^T \mathcal{E}(\mathbf{w}_t) - \mathcal{E}\left(\frac{1}{T} \sum_{t=1}^T \mathbf{e}_t^*\right) + \frac{\varepsilon}{T} \Omega_{\infty/2}((\boldsymbol{\Delta} \otimes \mathbf{I})\underline{\mathbf{e}}^*).$$

We can notice that the matrix $\mathbf{I} + \boldsymbol{\Delta}^\top \boldsymbol{\Delta} \in \mathbb{R}^{T \times T}$ is tridiagonal and for any $t = 2, \dots, T-1$ we have:

$$\mathbf{w}_t = (\mathbf{e}_t^* - \mathbf{e}_{t-1}^*) + (\mathbf{e}_t^* - \mathbf{e}_{t+1}^*),$$

while $\mathbf{w}_1 = \mathbf{e}_1^* - \mathbf{e}_2^*$ and $\mathbf{w}_T = \mathbf{e}_T^* - \mathbf{e}_{T-1}^*$. When applying \mathcal{E} and using sub-additivity, we therefore have for any $t = 2, \dots, T-1$:

$$\mathcal{E}(\mathbf{w}_t) \leq c(\mathcal{E}(\mathbf{e}_{t-1}^* - \mathbf{e}_t^*) + \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*))$$

with $\mathcal{E}(\mathbf{w}_1) \leq c\mathcal{E}(\mathbf{e}_1^* - \mathbf{e}_2^*)$ and $\mathcal{E}(\mathbf{w}_T) \leq c\mathcal{E}(\mathbf{e}_{T-1}^* - \mathbf{e}_T^*)$ where we used $c \geq 1$. Observing that if \mathcal{E} is symmetric around zero ($\mathcal{E}(\mathbf{z}) = \mathcal{E}(-\mathbf{z})$), the sum over $\mathcal{E}(\mathbf{w}_t)$ brings into play each term $\mathbf{e}_t^* - \mathbf{e}_{t-1}^*$ twice, we end up with

$$\frac{1}{T} \sum_{t=1}^T \mathcal{E}(\mathbf{w}_t) \leq 2c \frac{1}{T} \sum_{t=1}^{T-1} \mathcal{E}(\mathbf{e}_t^* - \mathbf{e}_{t+1}^*).$$

\square

5.4. Proofs of Lemma 4

Proof. We first apply the Cauchy-Schwartz inequality to get

$$\begin{aligned} |\hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)(\mathbf{x}_t - \hat{\mathbf{x}}_t)| &\leq \|\hat{\lambda}_t\|_2 \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|_2 \|\hat{\mathbf{A}}_t - \mathbf{A}_t\|_2 \\ &\leq 2R_\lambda R_x \|\hat{\mathbf{A}}_t - \mathbf{A}_t\|_F \end{aligned}$$

where we have used that the operator norm is smaller than the Frobenius norm. We then leverage Theorem 4 from (Hall and Willett, 2013): in their notation, we have $G_\ell \leq 1$ (since the sub-gradients of $\mathbf{A} \mapsto \|\mathbf{A}_t - \mathbf{A}\|_F$ are bounded by 1), $\sigma = 1$, $M \leq R_A/2$ and $D_{\max} \leq 2R_A^2$ along with $\Phi = \mathbf{I}$ and the learning rate $\eta_t = R_A/\sqrt{t}$. This leads to

$$\begin{aligned} \sum_{t=1}^T \|\hat{\mathbf{A}}_t - \mathbf{A}_t\|_F - \min_{\tilde{\mathbf{A}}_1, \dots, \tilde{\mathbf{A}}_T \in \mathcal{A}} \sum_{t=1}^T \|\tilde{\mathbf{A}}_t - \mathbf{A}_t\|_F \\ \leq \frac{2R_A^2}{\eta_{T+1}} + \frac{2R_A}{\eta_T} V + \frac{1}{2} \sum_{t=1}^T \eta_t \\ \leq \frac{2R_A}{\eta_T} (R_A + V) + R_A \sqrt{T} \\ \leq 3\sqrt{T}(V + R_A), \end{aligned}$$

where V stands for $\sum_{t=1}^T \|\mathbf{A}_t - \mathbf{A}_{t+1}\|_F$. \square

5.5. Proofs of Theorem 1

Proof. The proof consists in putting together the components introduced in the Section 3 of the core paper. We start from

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) = \frac{1}{T} \sum_{t=1}^T \mathcal{L}_t(\mathbf{x}_t^*, \lambda^*) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda})$$

which, using the decomposition (6), is equal to

$$\frac{1}{T} \sum_{t=1}^T (\mathcal{L}_t(\mathbf{x}_t^*, \lambda^*) - \mathcal{L}_t(\mathbf{x}_t^*, \hat{\lambda}_t)) + (\mathcal{L}_t(\mathbf{x}_t^*, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda})).$$

Moreover, Lemma A (see below) gives

$$\begin{aligned} \mathcal{L}_t(\mathbf{x}_t^*, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}) &\leq \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}) \\ &\quad + \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)(\mathbf{x}_t^* - \hat{\mathbf{x}}_t). \end{aligned}$$

Recalling the definitions of \mathcal{R}_T , \mathcal{G}_t and \mathcal{S}_A respectively in Lemma 1, (6) and (11), we obtain

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \leq \frac{1}{T} \sum_{t=1}^T \mathcal{G}_t(\hat{\lambda}_t) + \mathcal{R}_T + \mathcal{S}_A.$$

Noticing that with the choice $\varepsilon/T \triangleq G \sum_{t=1}^T \eta_t/T$, the relationship (8) holds, we finally obtain

$$\mathcal{P}^* - \mathcal{P}(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_T) \leq \mathcal{R}_e + \mathcal{R}_T + \mathcal{S}_A.$$

The rest of the proof follows by instantiating \mathcal{R}_T , \mathcal{S}_e and \mathcal{S}_A . The value of \mathcal{R}_T is given in Lemma 1, while the upper bound for \mathcal{S}_e is described in Lemma 3.

Moreover, ε/T is equal to $2R_\lambda G/\sqrt{T}$ and $GL \log(eT)/T$ in the convex and strongly convex cases respectively (given the η_t from Lemma 1). Finally, the upper bound on \mathcal{S}_A is given by Lemma 4. \square

6. Appendix: Technical lemmas

Lemma A. For any $\lambda \in \Lambda$ and $\mathbf{x}_t \in \mathcal{X}_t$, we have

$$\begin{aligned} \mathcal{L}_t(\mathbf{x}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda) &\leq \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda) \\ &\quad + \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)(\mathbf{x}_t - \hat{\mathbf{x}}_t). \end{aligned}$$

Proof. For any $\lambda \in \Lambda$ and $\mathbf{x}_t \in \mathcal{X}_t$, $\mathcal{L}_t(\mathbf{x}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda)$ is equal to

$$\begin{aligned} &= \hat{\mathcal{L}}_t(\mathbf{x}_t, \hat{\lambda}_t) + \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)\mathbf{x}_t - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda) \\ &\leq \hat{\mathcal{L}}_t(\hat{\mathbf{x}}_t, \hat{\lambda}_t) + \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)\mathbf{x}_t - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda) \\ &= \mathcal{L}_t(\hat{\mathbf{x}}_t, \hat{\lambda}_t) - \mathcal{L}_t(\hat{\mathbf{x}}_t, \lambda) + \hat{\lambda}_t^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)(\mathbf{x}_t - \hat{\mathbf{x}}_t), \end{aligned}$$

where we have used twice the relationship $\hat{\mathcal{L}}_t(\mathbf{x}_t, \lambda) = \mathcal{L}_t(\mathbf{x}_t, \lambda) - \lambda^\top (\hat{\mathbf{A}}_t - \mathbf{A}_t)\mathbf{x}_t$, and the fact that $\hat{\mathbf{x}}_t$ is defined as $\arg\max_{\mathbf{x}_t \in \mathcal{X}_t} \mathcal{L}_t(\mathbf{x}_t, \hat{\lambda}_t)$. \square

We next provide the details of the computations related to relationship between \mathcal{E} and \mathcal{E}^* presented in Table ???. We focus on one of the displayed instantiations, since the arguments in those other cases follow along the same lines.

Lemma B. For any $\mathbf{u} \in \mathbb{R}^m$,

$$\max_{\mathbf{w} \in \mathbb{R}^m} \left\{ \mathbf{u}^\top \mathbf{w} - \mathcal{I}_{\mathcal{B}_r \cap \mathbb{R}_+^m}(\mathbf{w}) - \frac{1}{2s} \|\mathbf{w}\|_2^2 \right\} = \mathcal{H}_{r,s}(\|\mathbf{u}\|_2),$$

where $\mathcal{H}_{r,s}(t) \triangleq \frac{1}{2} \min\{st^2, \frac{r^2}{s}\} + r[|t| - \frac{r}{s}]_+$.

Proof. We start by deriving the Lagrangian associated with

$$\begin{aligned} &\max_{\mathbf{w} \in \mathbb{R}^m} \left\{ \mathbf{u}^\top \mathbf{w} - \mathcal{I}_{\mathcal{B}_r \cap \mathbb{R}_+^m}(\mathbf{w}) - \frac{1}{2s} \|\mathbf{w}\|_2^2 \right\} \\ &= \max_{\mathbf{w} \in \mathcal{B}_r \cap \mathbb{R}_+^m} \left\{ \mathbf{u}^\top \mathbf{w} - \frac{1}{2s} \|\mathbf{w}\|_2^2 \right\} \end{aligned}$$

that is given by, for any $\alpha \geq 0, \beta \in \mathbb{R}_+^m$,

$$L(\mathbf{w}, \alpha, \beta) = \mathbf{u}^\top \mathbf{w} + \frac{\alpha}{2} (r^2 - \|\mathbf{w}\|_2^2) - \frac{1}{2s} \|\mathbf{w}\|_2^2 + \beta^\top \mathbf{w}.$$

Maximizing out \mathbf{w} , and recognizing the conjugate of the squared ℓ_2 norm (Example 3.27 in (Boyd and Vandenberghe, 2004)), we obtain the dual function

$$\max_{\mathbf{w} \in \mathbb{R}^m} L(\mathbf{w}, \alpha, \beta) = \frac{1}{2(\alpha + 1/s)} \|\mathbf{u} + \beta\|_2^2 + \frac{\alpha}{2} r^2.$$

In turn, we minimize with respect to the dual variables

$$\min_{\alpha \geq 0, \beta \in \mathbb{R}_+^m} \max_{\mathbf{w} \in \mathbb{R}^m} L(\mathbf{w}, \alpha, \beta) = \min_{\alpha \geq 0} \left\{ \frac{1}{2(\alpha + 1/s)} \|\mathbf{u}\|_2^2 + \frac{\alpha}{2} r^2 \right\}.$$

The optimal α can then be easily computed and is equal to

$$\left[\frac{\|\mathbf{u}\|_2}{r} - \frac{1}{s} \right]_+.$$

Plugging back this value into the dual function leads to the expression $\mathcal{H}_{r,s}(\|\mathbf{u}\|_2)$. The equality holds by invoking strong duality, which applies based on Slater's constraint qualification (Section 5.2.3 in (Boyd and Vandenberghe, 2004)). \square

Lemma C. For any $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, it holds that

$$\mathcal{H}_{r,s}(\|\mathbf{u} + \mathbf{v}\|_2) \leq 2 \left[\mathcal{H}_{r,s}(\|\mathbf{u}\|_2) + \mathcal{H}_{r,s}(\|\mathbf{v}\|_2) \right]$$

and

$$\mathcal{H}_{r,s}(\|\mathbf{u} + \mathbf{v}\|_2) \leq 2 \left[\mathcal{H}_{r,s}(\|\mathbf{u}\|_2) + \mathcal{H}_{r,s}(\|\mathbf{v}\|_2) \right]$$

where $\mathcal{H}_{r,s}(t) \triangleq \frac{1}{2} \min\{st^2, \frac{r^2}{s}\} + r[t - \frac{r}{s}]_+$.

Proof. Since $\mathcal{H}_{r,s}$ is non-decreasing and the triangle inequality holds for both $\|\cdot\|_2$ and $\|\cdot\|_+$, we concentrate on proving that for all $t_1, t_2 \in \mathbb{R}_+$,

$$\mathcal{H}_{r,s}(t_1 + t_2) \leq 2[\mathcal{H}_{r,s}(t_1) + \mathcal{H}_{r,s}(t_2)].$$

We analyze the 3 possible scenarios, namely (1) $t_1 + t_2 \leq \frac{r}{s}$ along with $t_1 \leq \frac{r}{s}$ and $t_2 \leq \frac{r}{s}$, (2) $t_1 + t_2 > \frac{r}{s}$ along with $t_1 \leq \frac{r}{s}$ and $t_2 > \frac{r}{s}$, and (3) $t_1 + t_2 > \frac{r}{s}$ along with $t_1 > \frac{r}{s}$ and $t_2 > \frac{r}{s}$.

Case (1) is handled by the fact that $(t_1 + t_2)^2 \leq 2(t_1^2 + t_2^2)$. As for case (2), we note that $\mathcal{H}_{r,s}(t_1 + t_2) \leq 2[\mathcal{H}_{r,s}(t_1) + \mathcal{H}_{r,s}(t_2)]$ is equivalent to

$$r(t_1 + t_2) - \frac{r^2}{2s} \leq st_1^2 + 2rt_2 - \frac{r^2}{s}$$

which can be rearranged as $\frac{r^2}{2s} \leq st_1^2 - rt_1 + rt_2$. The smallest possible value of the right-hand side is attained for $t_1 = \frac{r}{2s}$ and $t_2 = \frac{r}{s}$, thus leading to $\frac{r^2}{2s} \leq \frac{3r^2}{4s}$. Case (3) follows along the same line. \square

7. Appendix: Additional experiments for Section 4.1

We show below additional results when $\{\mathbf{A}_t, \mathbf{b}_t, \mathbf{u}_t\}$ are generated according to different distributions, namely,

Cauchy, uniform and gamma. As described in the protocol of Section 4.1, we continue to normalize $\{\mathbf{A}_t, \mathbf{b}_t, \mathbf{u}_t\}$ to unit norms. The same conclusions as those explained in Section 4.1 of the paper hold. We observe that some instantiations of \mathcal{E} and distribution, e.g., ℓ_1 with gamma in Figure 5, lead to settings where the additive relaxation (12) appears as tighter, although our non-additive approach still offers better reward versus constraint violation tradeoffs.

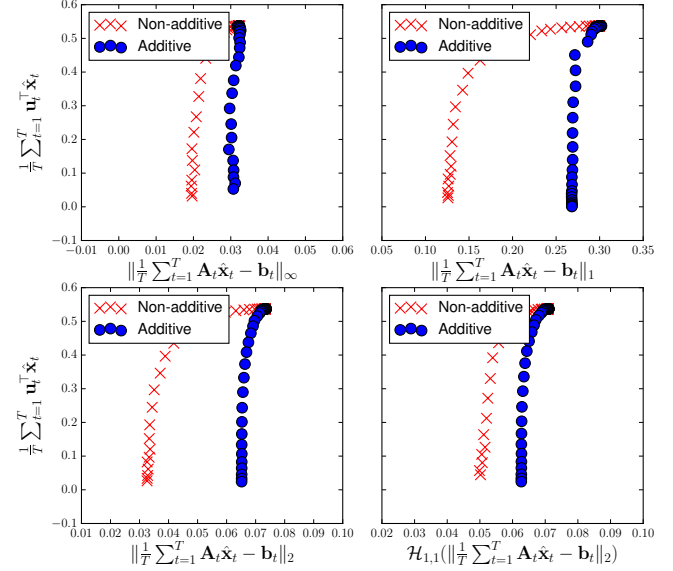


Figure 3. Reward $\frac{1}{T} \sum_{t=1}^T \mathbf{u}_t^\top \hat{\mathbf{x}}_t$ as a function of non-additive penalty $\mathcal{E}(\frac{1}{T} \sum_{t=1}^T \mathbf{A}_t \hat{\mathbf{x}}_t - \mathbf{b}_t)/R_\lambda$ for $R_\lambda = 2^\gamma$ with $\gamma \in \{-8, -7.5, \dots, 10\}$. The problem instances are generated with a standard Cauchy distribution. Red crosses correspond to our proposed online algorithm, while blue circles stand for a baseline algorithm with additive penalties. Each subplot displays a different instantiation of \mathcal{E} , namely $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_q$ for $q \in \{1, 2, \infty\}$ and $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \mathcal{H}_{1,1}(\|\mathbf{z}\|_2)$ where $\mathcal{H}_{1,1}$ is defined in Table 1.

References

- Agrawal, S. and Devanur, N. R. (2015). Fast algorithms for online stochastic convex programming. In *SODA 2015 (ACM-SIAM Symposium on Discrete Algorithms)*. SIAM-Society for Industrial and Applied Mathematics.
- Andrew, L., Barman, S., Ligett, K., Lin, M., Meyerson, A., Roytman, A., and Wierman, A. (2013). A tale of two metrics: simultaneous bounds on competitiveness and regret. In *ACM SIGMETRICS Performance Evaluation Review*, volume 41, pages 329–330.
- Awerbuch, B. and Kleinberg, R. (2008). Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114.
- Bera, S. K., Choudhury, A. R., Das, S., Roy, S., and




Figure 4. Reward $\frac{1}{T} \sum_{t=1}^T \mathbf{u}_t^\top \hat{\mathbf{x}}_t$ as a function of non-additive penalty $\mathcal{E}(\frac{1}{T} \sum_{t=1}^T \mathbf{A}_t \hat{\mathbf{x}}_t - \mathbf{b}_t)/R_\lambda$ for $R_\lambda = 2^\gamma$ with $\gamma \in \{-8, -7.5, \dots, 10\}$. The problem instances are generated with an uniform(-1,1) distribution. Red crosses correspond to our proposed online algorithm, while blue circles stand for a baseline algorithm with additive penalties. Each subplot displays a different instantiation of \mathcal{E} , namely $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_q$ for $q \in \{1, 2, \infty\}$ and $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \mathcal{H}_{1,1}(\|\mathbf{z}\|_2)$ where $\mathcal{H}_{1,1}$ is defined in Table 1.

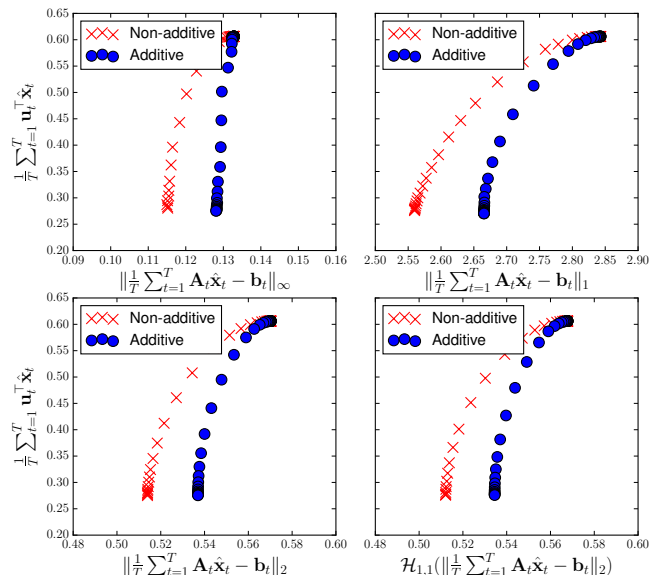


Figure 5. Reward $\frac{1}{T} \sum_{t=1}^T \mathbf{u}_t^\top \hat{\mathbf{x}}_t$ as a function of non-additive penalty $\mathcal{E}(\frac{1}{T} \sum_{t=1}^T \mathbf{A}_t \hat{\mathbf{x}}_t - \mathbf{b}_t)/R_\lambda$ for $R_\lambda = 2^\gamma$ with $\gamma \in \{-8, -7.5, \dots, 10\}$. The problem instances are generated with a gamma(2,2) distribution. Red crosses correspond to our proposed online algorithm, while blue circles stand for a baseline algorithm with additive penalties. Each subplot displays a different instantiation of \mathcal{E} , namely $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \|\mathbf{z}\|_q$ for $q \in \{1, 2, \infty\}$ and $\mathcal{E}(\mathbf{z}) = R_\lambda \cdot \mathcal{H}_{1,1}(\|\mathbf{z}\|_2)$ where $\mathcal{H}_{1,1}$ is defined in Table 1.

Thatchachar, J. S. (2013). Fenchel duals for drifting adversaries. Technical report, preprint arXiv:1309.5904.

Borodin, A. and El-Yaniv, R. (2005). *Online computation and competitive analysis*. Cambridge University Press.

Borodin, A., Linial, N., and Saks, M. E. (1992). An optimal on-line algorithm for metrical task system. *Journal of the ACM (JACM)*, 39(4):745–763.

Boyd, S. P. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Buchbinder, N., Chen, S., Naor, J., and Shamir, O. (2012). Unified algorithms for online learning and competitive analysis. In *Proceedings of the Annual Conference on Computational Learning Theory (COLT)*, pages 5–1.

Cesa-Bianchi, N., Gaillard, P., Lugosi, G., and Stoltz, G. (2012). A new look at shifting regret. Technical report, CoRR abs/1202.3323.

Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge University Press.

Chen, Y., Berkhin, P., Anderson, B., and Devanur, N. (2011). Online bidding algorithms for performance-based display ad allocation. In *Proceedings of the 17th*

ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2011), pages 1307–1315.

Diamond, S., Chu, E., and Boyd, S. (2014). CVXPY: A Python-embedded modeling language for convex optimization, version 0.2. <http://cvxpy.org/>.

Even-Dar, E., Kleinberg, R., Mannor, S., and Mansour, Y. (2009). Online learning for global cost functions. In *Proceedings of the annual conference on Computational Learning Theory (COLT)*.

Hall, E. and Willett, R. (2013). Dynamical models and tracking regret in online convex programming. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 579–587.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. Springer.

Hazan, E., Agarwal, A., and Kale, S. (2007). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192.

Helmbold, D., Schapire, R., Singer, Y., and Warmuth, M. (1996). On-line portfolio selection using multiplicative

- updates. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 243–251.
- Jadbabaie, A., Rakhlin, A., Shahrampour, S., and Sridharan, K. (2015). Online optimization: Competing with dynamic comparators. Technical report, preprint arXiv:1501.06225.
- Jaggi, M. (2013). Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 427–435.
- Kar, P., Narasimhan, H., and Jain, P. (2014). Online and stochastic gradient methods for non-decomposable loss functions. In *Advances in Neural Information Processing Systems*, pages 694–702.
- Koppel, A., Jakubiec, F. Y., and Ribeiro, A. (2014). A saddle point algorithm for networked online convex optimization. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pages 8292–8296. IEEE.
- Mahdavi, M., Jin, R., and Yang, T. (2012). Trading regret for efficiency: online convex optimization with long term constraints. *Journal of Machine Learning Research*, 13(1):2503–2528.
- Parikh, N. and Boyd, S. (2013). Proximal algorithms. *Foundations and Trends in optimization*, 1(3):123–231.
- Rakhlin, A., Sridharan, K., and Tewari, A. (2010). Online learning: Beyond regret. Technical report, preprint arXiv:1011.3168.
- Shalev-Shwartz, S. (2011). Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194.
- Shalev-Shwartz, S. and Kakade, S. M. (2009). Mind the duality gap: Logarithmic regret algorithms for online optimization. In *Advances in Neural Information Processing Systems*, pages 1457–1464.
- Shalev-Shwartz, S. and Singer, Y. (2006). Convex repeated games and Fenchel duality. In *Advances in Neural Information Processing Systems*, pages 1265–1272.
- Sion, M., of Scientific Research, O., States, U., and Force, A. (1957). *General Minimax Theorems*. United States Air Force, Office of Scientific Research.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the International Conference on Machine Learning (ICML)*.