
Between Stochastic and Adversarial Online Convex Optimization: Improved Regret Bounds via Smoothness

Sarah Sachs

University of Amsterdam
Korteweg-de Vries Institute for Mathematics
s.c.sachs@uva.nl

Hédi Hadiji

University of Amsterdam
Korteweg-de Vries Institute for Mathematics
hedi.hadiji@gmail.com

Tim van Erven

University of Amsterdam
Korteweg-de Vries Institute for Mathematics
tim@timvanerven.nl

Cristóbal Guzmán

Pontificia Universidad Católica de Chile
Institute for Mathematical and Computational Eng.
Facultad de Matemáticas and Escuela de Ingeniería
crguzmanp@mat.uc.cl

Abstract

Stochastic and adversarial data are two widely studied settings in online learning. But many optimization tasks are neither i.i.d. nor fully adversarial, which makes it of fundamental interest to get a better theoretical understanding of the world between these extremes. In this work, we establish novel regret bounds for online convex optimization in a setting that interpolates between stochastic i.i.d. and fully adversarial losses. By exploiting smoothness of the expected losses, these bounds replace dependence on the maximum gradient length by the variance of the gradients, which was previously known only for linear losses. In addition, they weaken the i.i.d. assumption by allowing, for example, adversarially poisoned rounds, which were previously considered in the expert and bandit setting. Our results extend this to the online convex optimization framework. In the fully i.i.d. case, our bounds match the rates one would expect from results in stochastic acceleration, and in the fully adversarial case, they gracefully deteriorate to match the minimax regret. We further provide lower bounds showing that our regret upper bounds are tight for all intermediate regimes in terms of the stochastic variance and the adversarial variation of the loss gradients.

1 Introduction

Two of the main approaches for solving convex optimization problems under uncertain data are stochastic convex optimization (SCO) [21, 30] and online convex optimization (OCO) [37]. These two models are very different in their assumptions and goals, despite the fact that they share many techniques. In SCO it is assumed that the loss functions follow an independent, identically distributed (i.i.d.) process, and the goal is to minimize the *excess risk*, which is the optimization error under the expected loss. By contrast, in OCO the losses can be chosen adversarially and the goal is to minimize

the *cumulative regret*, which is the difference between the cumulative incurred losses over rounds against the best fixed strategy in hindsight. Much less is known about what happens in between, in scenarios that interpolate between the i.i.d. and adversarial settings. This intermediate setting has drawn major attention in the recent years in the expert and bandit setting [14] [2] [36], however, as mentioned in [14], little is known for online convex optimization. Our work studies this in a generalization of the OCO setting, in which nature chooses distributions for the data that may vary arbitrarily over time, and we provide regret bounds in terms of two quantities that measure how adversarial these distributions are. The standard OCO setting corresponds to the case where the distributions are point-masses on adversarial data points.

Main Contribution. Our main contribution is a new analysis of *optimistic online algorithms* [25, 26] that takes advantage of smoothness of the expected loss. This analysis allows for a gradual interpolation between worst-case adversarial regret bounds and the best-known expected regret bounds in the stochastic case, and also provides quantifiable improvements for intermediate cases.¹ To capture the full range between i.i.d. and fully adversarial settings, we consider a similar adversarial model as in [27], i.e., nature chooses distributions \mathcal{D}_t in iteration t , and the learner suffers loss $f_t(x_t, \xi_t)$ with $\xi_t \sim \mathcal{D}_t$. Importantly, we do not assume that the distributions \mathcal{D}_t are all the same, but they may vary adversarially over time.

To properly quantify the interpolation between the i.i.d. and fully adversarial settings in the regret bound, we introduce two parameters for the loss sequence. Namely the *cumulative variance*, $\bar{\sigma}_T$, which captures the stochastic aspect of the learning task, i.e., the variance of the \mathcal{D}_t ; and the *cumulative adversarial variation*, $\bar{\Sigma}_T$, which captures the adversarial difficulties of the data, i.e., the difference between $\mathbb{E}_{\xi \sim \mathcal{D}_t}[\nabla f(\cdot, \xi)]$ and $\mathbb{E}_{\xi \sim \mathcal{D}_{t-1}}[\nabla f(\cdot, \xi)]$. With these two key quantities, our first main result in Theorem 5 shows that the expected regret $\mathbb{E}[R_T(u)]$, that is the difference between the cumulative losses of the learner and a fixed solution in hindsight, is bounded by

$$\mathbb{E}[R_T(u)] = O\left(D(\bar{\sigma}_T + \bar{\Sigma}_T)\sqrt{T} + LD^2\right), \quad (1)$$

where L is the smoothness constant of the expected functions $F^t = \mathbb{E}_{\xi \sim \mathcal{D}_t}[f(\cdot, \xi)]$. If, in addition, the functions F^t are μ -strongly convex, then in Theorem 7 we obtain (below, $\kappa := L/\mu$ is the condition number)

$$\mathbb{E}[R_T(u)] = O\left(\frac{1}{\mu}(\sigma_{\max}^2 + \Sigma_{\max}^2) \log T + LD^2 \kappa \log \kappa\right).$$

Both bounds are tight: we prove matching lower bounds in Theorems 6 and 8. In Section 3.1 we show that our results match the known adversarial regret bounds as well as the best results in the i.i.d. case. For the latter, only the linear case was so far obtained directly via a regret analysis (see Sec. 5.2 in [25], and prior work [12, 4]). Using optimistic mirror descent, they obtained the regret guarantee of $R_T(u) \leq O(\sqrt{\sum_{t=1}^T \|\nabla f(x_t, \xi_t) - m_t\|^2})$, where m_t denotes an optimistic guess of the gradient that is chosen before round t . In the i.i.d. case with the prediction $m_t = \nabla f(x_{t-1}, \xi_{t-1})$, this can be shown to imply that the expected regret is upper bounded by $\mathbb{E}[R_T(u)] = O(\sigma\sqrt{T} + \sqrt{\sum_{t=1}^T \mathbb{E}[\|\nabla F^t(x_t) - \nabla F^t(x_{t-1})\|^2]})$, where σ denotes the variance of the stochastic gradients. This simplifies to

$$\mathbb{E}[R_T(u)] = O(\sigma\sqrt{T}) \quad \text{for the i.i.d. case with linear functions } \nabla F^t,$$

which is a special case of (1), because $\sigma = \bar{\sigma}_T$ and $\bar{\Sigma}_T = 0$ for i.i.d. losses, and $\nabla F^t(x_t) = \nabla F^t(x_{t-1})$ and $L = 0$ for linear functions. It is not immediately obvious how to generalize this result to general convex functions with smoothness $L > 0$, however. In this case, we can guess the appropriate regret bound based on known convergence results for *stochastic accelerated gradient descent* (SAGD) [17, 8]: if we knew in advance that the losses would be i.i.d. and we did not care about computational efficiency, then we could run a new instance of SAGD for each round t . Summing the known rate for SAGD over t then gives $\mathbb{E}[R_T(u)] \leq O(\sigma\sqrt{T} + L)$ (for more details see the batch-to-online conversion in Appendix D). This raises the question if a similar bound can also be obtained directly via a regret analysis, without assuming i.i.d. observations in advance. This

¹It is well-known that in the fully adversarial case smoothness does not yield asymptotic improvements on regret [11], whereas for SCO improvements can be obtained only under low-noise [8].

question is then answered by our result [1], which indeed reduces to this rate for general convex smooth i.i.d. functions, matching the aforementioned bound up to constants. We achieve this by using smoothness to bound $\mathbb{E}[\|\nabla F^t(x_t) - \nabla F^t(x_{t-1})\|^2] \leq L \mathbb{E}[\|x_t - x_{t-1}\|^2]$, which can be canceled by a negative quadratic term that we obtain from an improved analysis of the regret. The use of this negative term in the analysis dates back to [20], who used it to achieve an improved $O(1/T)$ rate on the extra-gradient method.

In addition to unifying the analysis of these two extreme cases and obtaining the best-known results via one algorithm (OFTRL [2] for convex functions and OFTRL on a surrogate loss [5] for strongly convex functions), our results give new insight for intermediate cases. Thus, as a second main contribution we shed light on a setting that is neither fully adversarial nor i.i.d. To illustrate this, we highlight some examples here, which received attention in the recent literature.

Adversarial corruptions: Consider i.i.d. functions with adversarial corruptions, as considered in the context of the expert and bandit settings in [14], [2]. If the (cumulative) corruption level is bounded by a constant C , in [14] an expected regret bound of $\mathbb{E}[R_T(u)] = O(R_T^s + \sqrt{CR_T^s})$ was obtained, where R_T^s denotes the regret with respect to the uncorrupted data. In [14], the authors raised the question of whether it is possible to obtain regret bounds with a similar square-root dependence on the corruption level C for online convex optimization. Indeed, for this intermediate model, we derive a regret bound

$$\mathbb{E}[R_T(u)] \leq O(R_T^s + D\sqrt{GC}),$$

for the general convex case from our Theorem 5. We elaborate on this in Section 4.2.

Random order models The random order model (ROM) dates back to [18] in combinatorial online learning. It has drawn attention in the online convex optimization community as an elegant relaxation of the adversarial model [31, 6]. Complementary to the results in [31], we show that the dependence on G in the regret bound can be reduced to a dependence of σ , where σ denotes the variance of gradients in the uniform distribution over loss functions f_1, \dots, f_T . That is,

$$\mathbb{E}[R_T(u)] \leq O\left(D\sigma\sqrt{T\log\left(e\frac{\tilde{\sigma}}{\sigma}\right)}\right),$$

where $\tilde{\sigma}$ denotes a slightly weaker notion of variance (see Corollary [10]). We derive these results from our main theorem under stronger assumptions than in [31], namely, we require individual convexity of the adversarial choices whereas [31] only assume cumulative convexity. However, we also obtain a better rate with σ instead of G as the leading factor, so the results are not directly comparable. We also consider a variant of the random order model, which we call the *multiple pass random order model (multi-pass ROM)*. This is inspired by multiple shuffle SGD and can be considered another intermediate example between adversarial and stochastic data. We elaborate on both examples, i.e., the ROM and multi-pass ROM in Section 4.3.

1.1 Related work

As mentioned in the previous section, our work is inspired by results in the *gradual variation* and in the *stochastic approximation* literature. The gradual variation literature dates back to [12], with later extensions by [4] and [25, 26]. In addition to some technical relation to the aforementioned work, there is also a natural relation between our parameters $\bar{\sigma}_T$ and $\bar{\Sigma}_T$ and variational parameters in [33], [12] or [4]. However, as we elaborate in Remark 4, there are some fundamental differences between these variational parameters and $\bar{\sigma}_T, \bar{\Sigma}_T$, which prevent us from directly obtaining a smooth interpolation from these results. It is also interesting to note that there is a relation between $\bar{\Sigma}_T$ and the *path length* parameters considered in *dynamic regret* bounds [35, 34]. However, since their analysis targets a fundamentally different notion of regret, namely dynamic regret, the results are incomparable.

With respect to the results, our findings are fundamentally different from the stochastic approximation literature, since we do not rely on the assumption that the data is following a distribution. However, we were inspired by analysis techniques and the convergence thresholds set by this literature. Our approach of obtaining accelerated rates by negative terms arising from smoothness in a regret bound has previously been used in the context of variational inequalities and saddle-point problems. Using

this idea, [20] obtained improved rates $O(1/T)$ for the extra-gradient method. More recently [17] showed that acceleration in stochastic convex optimization can benefit from negative terms arising in optimistic FTRL via an anytime-online-to-batch conversion [5]. Although an important inspiration for our approach, the techniques of [17] do not directly carry over, because they evaluate gradients at the time-average of the algorithm’s iterates, making them much more stable than the last iterate, which comes up when controlling the regret. Algorithms used both in SCO and OCO follow a vast literature on stochastic approximation methods, e.g. [28, 21, 24]. For this work, we are particularly interested in the more recent literature on acceleration in SCO [8, 15, 17]. In this research field efficiency is traditionally measured in terms of excess risk. On the one hand, regret upper bounds can be converted into excess risk bounds, through the so-called *online-to-batch conversions* [3]; on the other hand, excess risk guarantees do not directly lead to regret bounds, and even if they do some key features of the rates might be lost. These latter methods, known as *batch-to-online conversions* are discussed in Appendix D.

Outline In Section 2 after setting up notation and basic definitions, we introduce the *stochastically extended adversarial* model, a generalization of the standard adversarial model similar to the model used in *smoothed analysis*. Our main results can be found in Section 3. In Section 4 we illustrate our results by highlighting several special cases, such as the random order model and the adversarially corrupted stochastic model. Finally, in Section 5 we set our findings into a broader context and give perspective for future work.

2 Setting

We recall the *online convex optimization (OCO)* problem. Here, we consider a sequence of convex functions f_1, \dots, f_T defined over a closed and bounded convex set $\mathcal{X} \subseteq \mathbb{R}^d$, which become available to the learner sequentially. In the standard *adversarial model*, the learner chooses $x_t \in \mathcal{X}$ in round t , then function f_t is revealed and the learner suffers loss $f_t(x_t)$. The success of the learner is measured against all fixed $u \in \mathcal{X}$. Hence, the goal of the learner is to minimize the *regret*, that is, the difference between their cumulative loss $\sum_{t=1}^T f_t(x_t)$ and that of the best fixed choice in hindsight, namely $\min_{u \in \mathcal{X}} \sum_{t=1}^T f_t(u)$.

Throughout the paper we use the notation $[T] = \{1, \dots, T\}$. We follow the convention that δ_c denotes a Dirac measure at a point c , and $\|\cdot\|$ denotes the Euclidean norm.

2.1 Stochastically Extended Adversarial Model

We extend the aforementioned adversarial model by letting nature choose a distribution \mathcal{D}_t from a set of distributions. Then the learner suffers loss $f(x_t, \xi_t)$ where $\xi_t \sim \mathcal{D}_t$. Note that if the set of distributions is sufficiently rich, this model contains the standard adversarial model and the stochastic model as special cases (see Examples 1, 2).

We introduce some notation to make this more precise. Let $\mathcal{X} \subset \mathbb{R}^d$ be a closed convex set and Ξ a measurable space. Define $f : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$ and assume $f(\cdot, \xi)$ is convex over $\xi \in \Xi$. Suppose \mathcal{D} is a set of probability distributions over Ξ . For any $\mathcal{D} \in \mathcal{D}$, we denote the gradient mean by $\nabla F^{\mathcal{D}}(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [\nabla f(x, \xi)]$ and the function mean by $F^{\mathcal{D}}(x) := \mathbb{E}_{\xi \sim \mathcal{D}} [f(x, \xi)]$. Furthermore, denote by $\sigma_{\mathcal{D}}^2$ an upper bound on the variance of the gradients

$$\sigma_{\mathcal{D}}^2 = \max_{x \in \mathcal{X}} \mathbb{E}_{\xi \sim \mathcal{D}} \left[\|\nabla f(x, \xi) - \nabla F^{\mathcal{D}}(x)\|^2 \right].$$

We introduce some shorthand notation when distributions are indexed by rounds. Given $t \in [T]$, we write F^t and σ_t^2 instead of $F^{\mathcal{D}_t}(x)$ and $\sigma_{\mathcal{D}_t}^2$, respectively. Let us now introduce the *stochastically extended adversary protocol*.

Definition 1 (Stochastically Extended Adversary (SEA)). *In each round t , the learner chooses $x_t \in \mathcal{X}$, the SEA picks $\mathcal{D}_t \in \mathcal{D}$. The learner and the SEA then both observe a sample $\xi_t \sim \mathcal{D}_t$, and the learner suffers loss $f(x_t, \xi_t)$.*

We note that the SEA model is closely related to the adversarial model considered in the context of *smoothed analysis* [27, 10, 32]. However, in contrast to this line of work, we do not focus on

distributions with sufficient anti-concentration (c.f., Def. 1.1 in [10]). Indeed, this restriction would exclude, among others, the fully adversarial case as described below. Note also that we assume that the SEA is an adaptive adversary, which has access to the realizations ξ_t , so the distribution \mathcal{D}_{t+1} may depend on ξ_1, \dots, ξ_t . This assumption is not relevant for the fully adversarial or the i.i.d. settings. In the former, because there is no randomness, and in the latter, because there is no change in distribution. However, it is important for some of the intermediate cases, and in particular for the random order model. The SEA model contains several common settings from the literature as special cases. To illustrate this, we list some examples.

1. **Adversarial Model:** The SEA chooses a Dirac measure $\delta_{c_t} \in \mathfrak{D}$ in each round. Then for any $\xi_t \sim \delta_{c_t}$, the SEA selects $f(\cdot, \xi_t)$, and the model reduces to an adversary selecting directly the functions $f_t(\cdot)$.
2. **Stochastic I.I.D. Model:** The SEA chooses a fixed $\mathcal{D} \in \mathfrak{D}$ and selects $\mathcal{D}_t = \mathcal{D}$ at each round t .
3. **Adversarially Corrupted I.I.D. Model:** The adversary selects an i.i.d. source \mathcal{D} and perturbs the data with adversarial corruptions. This fits in our framework as follows: given a corruption level $C \geq 0$, the SEA chooses distributions $\mathcal{D}_t = \mathcal{D} \otimes \delta_{c_t}$ such that $\sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\mathbb{E}_{\xi \sim \mathcal{D}}[\nabla f(x, \xi)] - \mathbb{E}_{\xi' \sim \mathcal{D}_t}[\nabla f(x, \xi')]\| \leq C$.
4. **Random Order Model (ROM):** Among a fixed family of losses $\mathcal{F} = (f_i, i \in [n])$, the SEA randomly picks functions in \mathcal{F} via sampling without replacement, possibly performing multiple passes over the losses and reshuffling between the passes. Formally, define $\Xi = [n]$, and $f(x, \xi) = f_\xi$ for $\xi \in \Xi$. If $t \in [nk, n(k+1)]$ is in epoch $k \in \mathbb{N}$, then the SEA chooses the distribution $\mathcal{D}_t = \text{Unif}(\Xi \setminus \{\xi_s : s \in [nk+1, n(t-1)]\})$, which is uniform on the functions that have not yet been sampled in epoch k .

To quantify the hardness of the loss sequence, we introduce the *cumulative stochastic variance* and *adversarial variation*; we also define an average of these quantities.

We denote by \mathbb{E} the expectation taken with respect to the joint distribution of $(x_1, \xi_1, \dots, x_T, \xi_T)$. Note that the choice of the adversary \mathcal{D}_t can be random itself, as it depends on the past observations (of both the player's actions and the realizations of the ξ_t 's). In this case, σ_t is also a random quantity.

Definition 2 (Cumulative Stochastic Variance and Cumulative Adversarial Variation). *Suppose the SEA chooses distributions $\mathcal{D}_1, \dots, \mathcal{D}_T$. Recall that σ_t^2 is a shorthand for $\sigma_{\mathcal{D}_t}^2$. The cumulative stochastic variance and the cumulative adversarial variation are defined as*

$$\sigma_{[1:T]}^{(2)} = \mathbb{E} \left[\sum_{t=1}^T \sigma_t^2 \right] \quad \text{and} \quad \Sigma_{[1:T]}^{(2)} = \mathbb{E} \left[\sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\nabla F^t(x) - \nabla F^{t-1}(x)\|^2 \right].$$

We also let $\bar{\sigma}_T$ and $\bar{\Sigma}_T$ denote the square root of the average stochastic variance or adversarial variation, respectively; that is, $\bar{\sigma}_T^2 = \sigma_{[1:T]}^{(2)}/T$ and $\bar{\Sigma}_T^2 = \Sigma_{[1:T]}^{(2)}/T$.

Note that in the special case when all f_t are fully adversarial, $\bar{\sigma}_T = 0$. On the contrary, in the stochastic case, i.e., if all for each round t , the distribution \mathcal{D}_t is equal to a fixed (but arbitrarily chosen) \mathcal{D} , then $\bar{\Sigma}_T = 0$. In this case, $\bar{\sigma}_T$ reduces to the common definition of the gradient variance upper bound in the SCO literature [7, 8]. If however, the SEA chooses one distribution \mathcal{D}_i for the first rounds and then switches to a different distribution \mathcal{D}_j , then $\bar{\sigma}_T$ can only be upper bounded by $\max(\sigma_i, \sigma_j)$. This upper bound can be pessimistic, however, for some results it gives a better intuition. For this purpose we also define the *maximal stochastic variance* and *maximal adversarial variation*.

Definition 3 (Maximal Stochastic Variance and Maximal Adversarial Variation). *Let σ_{\max}^2 be an upper bound on all variances σ_t^2 . That is,*

$$\sigma_{\max}^2 = \max_{t \in [T]} \mathbb{E}[\sigma_t^2] \quad \text{and} \quad \Sigma_{\max}^2 = \max_{t \in [T]} \mathbb{E} \left[\sup_{x \in \mathcal{X}} \|\nabla F^t(x) - \nabla F^{t-1}(x)\|^2 \right].$$

Remark 4. *As we mentioned in the introduction, the cumulative stochastic variance and the adversarial variation have some similarities with parameters in gradual variation regret bounds. For linear functions $\langle \mu_t, \cdot \rangle$, the bounds in [12] involve the parameter $\text{Var}_T = \sum_{t=1}^T \|\mu_t - \bar{\mu}_T\|^2$ where $\bar{\mu}_T$ is*

the average of the gradients. For OCO with general convex functions, [4] provide upper bounds on the regret in terms of the L_p -deviation $D_p = \sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\nabla f_t(x) - \nabla f_{t-1}(x)\|_p^2$. In Lemmas I3 and I4 in Appendix A we show that in the SEA framework, both of these types of bounds are generally worse than ours, and that the difference can be arbitrarily large.

In [26] the regret is bounded in terms of $\sum_{t=1}^T \|g_t - M_t\|^2$. As mentioned in the introduction, unless the loss functions are linear or the learner has knowledge of the gradient mean, $\sum_{t=1}^T \|g_t - M_t\|^2$ cannot directly be reduced to $\sigma_{[1:T]}^{(2)}$ or $\Sigma_{[1:T]}^{(2)}$.

2.2 Assumptions

In our analysis, we will frequently use several of the following additional assumptions. Some of these were already mentioned in the introduction. We keep them all together here, for the convenience of the reader and for clear reference. For any $\mathcal{D} \in \mathfrak{D}$:

- (A0) the adversary has access to independent samples $\xi \sim \mathcal{D}$.
- (A1) the function $f(\cdot, \xi)$ is convex, and gradients are bounded by G a.s. when $\xi \sim \mathcal{D}$.
- (A2) the expected function $F^{\mathcal{D}}$ is L -smooth, i.e. $\nabla F^{\mathcal{D}}$ is L -Lipschitz continuous.
- (A3) for any $x \in \mathcal{X}$, the variance $\mathbb{E}_{\xi \sim \mathcal{D}}[\|\nabla f(x, \xi) - \nabla F^{\mathcal{D}}(x)\|^2]$ is finite.
- (A4) the expected function $F^{\mathcal{D}}(\cdot)$ is μ -strongly convex.

We assume that (A0) always holds. Assumptions (A1) (A2) and (A3) are standard in stochastic optimization, and are similar to common assumptions for online convex optimization. There, it is typically assumed that the adversarial samples $f_t(\cdot)$ are convex (or even linear) and the gradient norms $\|\nabla f_t(\cdot)\|$ are bounded. Note that we only require gradient Lipschitz continuity and strong convexity to hold for the expected loss.

3 Algorithms and Regret Bounds

3.1 Convex Smooth Functions

We use *Optimistic Follow-the-Regularised-Leader* (OFTRL) (see, e.g., [16, 25]) to minimize regret. Let $(\eta_t)_{t \in [T]}$ be a non-decreasing and positive sequence of stepsizes, possibly tuned adaptively with the observations. At each step t , the learner makes an optimistic prediction $M_t \in \mathbb{R}^d$ and updates its iterates as

$$x_t = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \left\langle x, M_t + \sum_{s=1}^{t-1} g_s \right\rangle + \frac{\|x\|^2}{\eta_t} \right\}, \quad (2)$$

where we denoted by $g_t = \nabla f(x_t, \xi_t)$ the observed gradient at time t . To state our results, we denote by $\mathbb{E}[\cdot]$ the expectation with respect to the joint distribution of $(x_1, \xi_1, \dots, x_T, \xi_T)$. Our objective is to bound the average regret:

$$\mathbb{E}[R_T(u)] := \mathbb{E} \left[\sum_{t=1}^T \langle g_t, x_t - u \rangle \right].$$

The following theorem, proved in Appendix B.1 is our main result for convex functions.

Theorem 5. Fix a user-specified parameter $\nu > 0$. Under assumptions (A1) (A2) (A3) OFTRL, with $M_t = g_{t-1}$ and adaptive step-size $\eta_t = D^2 / (\nu + \sum_{s=1}^{t-1} \eta_s \|g_s - M_s\|^2)$, has regret

$$\mathbb{E}[R_T(u)] \leq D(6\bar{\sigma}_T + 3\sqrt{2}\bar{\Sigma}_T)\sqrt{T} + \frac{3\sqrt{2}DG}{2} + \nu + \frac{1}{\nu}(4D^2G^2 + 9L^2D^4). \quad (3)$$

The algorithm needs to know only D . If G and L are also known, one can tune $\nu = LD^2 + D^2G^2$ to get

$$\mathbb{E}[R_T(u)] \leq O(D(\bar{\sigma}_T + \bar{\Sigma}_T)\sqrt{T} + DG + LD^2).$$

Moreover, if only convexity of the individual losses holds (A1) then tuning $\nu = 2DG$ ensures the (deterministic) bound $R_T(u) \leq 3\sqrt{2}DG\sqrt{T} + 4DG$.

Without prior knowledge of the smoothness parameter, the best the player can do is to tune ν according to a guessed value L_0 . This affects the constants in the bound by an additive term of order $(L_0 + L^2/L_0)D^2$; it would be interesting to determine if this is an inevitable price to pay for the lack of knowledge of L . A similar discussion can be held for G . Note that the worst-case regret bound of order $DG\sqrt{T}$ always holds every time OFTRL is used in this article, even without smoothness. To avoid distraction, we will not recall this fact in the applications.

The algorithm and analysis dwell on two ideas: the adaptive tuning of the learning rate à la Ada-Hedge/AdaFTRL [19, 23] with optimism, together with the fact that we keep a negative Bregman divergence term in the analysis, which is crucial to obtain our bound.

Lower Bound. The upper bound in Theorem 5 is tight up to additive constants, as the following result shows.

Theorem 6. *For any learning algorithm, and for any pair of positive numbers (σ, Σ) there exists a function $f : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$ and a sequence of distributions satisfying assumptions (A1) (A2) (A3) with $\bar{\sigma}_T \geq \sigma$ and $\bar{\Sigma}_T \geq \Sigma$ such that*

$$\mathbb{E}[R_T(u)] \geq \Omega(D(\bar{\sigma}_T + \bar{\Sigma}_T)\sqrt{T}).$$

The proof, in Appendix B.2 relies on a lower bound from stochastic optimization [1, 21] together with the fact that we can construct a sequence of convex and L -smooth loss functions such that $\bar{\Sigma}_T$ is in the order of the gradient norms G . Combining these insights with the lower bound $\Omega(DG\sqrt{T})$ [23] gives the desired result.

3.2 Strongly Convex and Smooth Functions

Up to this point, we have only considered functions which satisfy the weaker set of assumptions (A1) (A2) (A3). In this section, we show what improvements can be achieved if strong convexity also holds, that is, if (A4) is satisfied with some known parameter $\mu > 0$. For $g_t = \nabla f(x_t, \xi_t)$, define the surrogate loss function

$$\ell_t(x) = \langle g_t, x - x_t \rangle + \frac{\mu}{2} \|x - x_t\|^2. \quad (4)$$

We use Optimistic Follow-the-Leader (OFTL) on the surrogate losses. For each step t , the learner makes an optimistic prediction of the next gradient $M_t \in \mathbb{R}^d$ and selects

$$x_t = \operatorname{argmin}_{x \in \mathcal{X}} \left\{ \sum_{s=1}^{t-1} \ell_s(x) + \langle M_t, x \rangle \right\}. \quad (5)$$

The next theorem is analogous to Theorem 5 for curved losses, and will be our main tool in establishing results for strongly convex losses; see Appendix B.3 for a proof.

Theorem 7. *Under assumptions (A1) (A4), the expected regret of OFTL with $M_t = \nabla f(x_{t-1}, \xi_{t-1})$ on surrogate loss functions ℓ_t defined in (4) is bounded as*

$$\begin{aligned} \mathbb{E}[R_T(u)] &\leq \frac{1}{\mu} \sum_{t=1}^T \frac{1}{t} \left(8\sigma_{\max}^2 + 4\mathbb{E} \left[\sup_{x \in \mathcal{X}} \|\nabla F^t(x) - \nabla F^{t-1}(x)\|^2 \right] \right) + \frac{4D^2L^2}{\mu} \log \left(1 + \frac{16L}{\mu} \right) \\ &\leq \frac{1}{\mu} (8\sigma_{\max}^2 + 4\Sigma_{\max}^2) \log T + \frac{4D^2L^2}{\mu} \log \left(1 + \frac{16L}{\mu} \right). \end{aligned}$$

Note that OFTL requires no tuning besides the strong convexity parameter used in the surrogate losses. In particular, it is adaptive to the smoothness L .

Lower Bound The bound in Theorem 7 is tight, as the next result, proved in Appendix B.4 shows.

Theorem 8. *For any learning algorithm, and for any pair of positive numbers (σ, Σ) there exists a function $f : \mathcal{X} \times \Xi \rightarrow \mathbb{R}$ and a sequence of distributions satisfying assumptions (A1) (A2) (A3) and (A4) with $\sigma_{\max} \geq \sigma$ and $\Sigma_{\max} \geq \Sigma$ such that*

$$\mathbb{E}[R_T(u)] \geq \Omega \left(\frac{1}{\mu} (\sigma_{\max}^2 + \Sigma_{\max}^2) \log T \right).$$

4 Implications

We derive consequences of our results from Section 3. Further examples can be found in Appendix E.

4.1 Interpolating Known Results: Fully Adversarial and i.i.d. Data

A first implication of our analysis is that we recover both the adversarial and i.i.d. rates, *via a single adaptive algorithm*.

Convex Case For adversarial data, $\sigma_t = 0$ for all t , and $\Sigma_{[1:T]}^2 \leq G\sqrt{2T}$. Thus, Theorem 5 guarantees a bound $R_T(u) \leq O(DG\sqrt{T})$, which is known to be the optimal rate up to the additive constants, cf. [37] (note that the expectation does not act on the regret in this case). Simultaneously, if the data is i.i.d., then Theorem 5 guarantees that

$$\mathbb{E}[R_T(u)] \leq O(D\sigma\sqrt{T} + LD^2 + DG). \quad (6)$$

From standard online-to-batch conversion, this implies an excess risk for the related SCO problem of order $O(D\sigma/\sqrt{T} + D(LD + G)/T)$, which matches the well-known result by [7] up to lower order terms. On the other hand, using batch-to-online conversion (see Appendix D) with the best known accelerated convergence result in SCO, gives $O(D\sigma\sqrt{T} + LD^2)$ regret. Therefore, up to a constant, our result coincides with the best-known results from SCO. Note that also generalizes the improvement obtained for linear functions in the i.i.d. setting in [25, Section 6.2].

Strongly Convex Case The adaptive interpolation between i.i.d. and adversarial rates also holds in the strongly convex case. For adversarial data, the bound of Theorem 7 is of order $(G^2/\mu) \log T$, which is known to be the optimal worst-case rate, cf. [13]. For i.i.d. data, the dependence on G^2 improves to σ^2 , yielding a bound of order $O((\sigma^2/\mu) \log T + LD^2 \kappa \log \kappa)$. This improvement is akin to improvements obtained by *accelerated stochastic gradient descent* in the context of stochastic optimization [8, 17]. In fact, applying batch-to-online conversions and summing the optimization rates would yield a regret bound similar to ours; c.f. Appendix D.

4.2 Adversarially Corrupted Stochastic Data

We consider a natural generalization to online convex optimization of the corruption model considered in the bandit literature [29, 36], also recently studied in [14] for prediction with expert advice. There, the author obtains a regret bound that is the sum of the i.i.d. rate and of a term of order \sqrt{C} where C is the total amount of perturbation. They then raise the open question of whether similar results could be obtained for general convex losses. We provide a positive answer to this question in this section, with the regret bound in Corollary 9.

In this model, the generating process of the losses is decomposed as a combination of losses coming from i.i.d. data, with a small additive adversarial perturbation. This fits in the framework by setting $\xi_t = (\xi_{\text{iid},t}, c_t) \sim \mathcal{D}_t = \mathcal{D} \otimes \delta_{c_t}$ and

$$f(x, \xi_t) = h(x, \xi_{\text{iid},t}) + c_t(x)$$

where c_t is the adversarial part of the losses selected by the adversary, and $\xi_{\text{iid},t} \sim \mathcal{D}$ is a sequence of identically distributed random variables. Note that, similarly to our inspirations [14, 29], and contrary to other corruption models for prediction with expert advice [2], we measure the regret against the perturbed data. Define $F = \mathbb{E}_{\xi \sim \mathcal{D}}[h(\cdot, \xi)]$, so that $F^t(x) = F(x) + c_t(x)$. The amount of perturbation is measured by a parameter $C > 0$ bounding

$$\sum_{t=1}^T \max_{x \in \mathcal{X}} \|\nabla c_t(x)\| \leq C,$$

which is a natural measure of perturbation on the feedback used by the player (note that adding a constant to the perturbations does not change the regret). In this case, the adversarial perturbation on the loss does not affect the variance and $\sigma_{\mathcal{D}_t}^2 = \sigma_{\mathcal{D}}^2$. The perturbation appears in the loss variation as for any $t \geq 2$, for any $x \in \mathcal{X}$,

$$\begin{aligned} \|\nabla F^t(x) - \nabla F^{t-1}(x)\|^2 &\leq 2G\|\nabla F^t(x) - \nabla F^{t-1}(x)\| \\ &= 2G\|\nabla c_t(x) - \nabla c_{t-1}(x)\| \leq 2G(\|\nabla c_t(x)\| + \|\nabla c_{t-1}(x)\|). \end{aligned}$$

Upon taking the supremum over $x \in \mathcal{X}$ and summing over t , we get (with the convention that $c_0 \equiv 0$),

$$\Sigma_{[1:T]}^{(2)} = \sum_{t=1}^T \sup_{x \in \mathcal{X}} \|\nabla c_t(x) - \nabla c_{t-1}(x)\|^2 \leq 4GC.$$

Hence, Theorem 5 combined with the bounds on $\bar{\sigma}_T$ and $\bar{\Sigma}_T$ yields the following regret guarantee.

Corollary 9. *In the adversarially corrupted stochastic model, adaptive OFTRL enjoys the bound*

$$\mathbb{E}[R_T(u)] = O(D\sigma\sqrt{T} + D\sqrt{GC}).$$

This regret bound is the sum of the i.i.d. rate for the unperturbed source with a term sublinear in the amount of perturbations C , achieved without prior knowledge of C . This provides an answer to a question by [14]. The construction of the lower bound Theorem 6 also provides a partial lower bound for the corruption model, proving some worst-case optimality of our upper bound, see Appendix B.2.1.

An interesting open question that remains would be to extend these results to strongly convex losses.

Regret against unperturbed losses Depending on modeling choices, one might consider the perturbations c_t to be a true part of the data that deviate from an i.i.d. model, or one may see them as adversarial noise which should be discarded. In the latter case, regret against the i.i.d. part of the data is most relevant. As noticed by [2, Observation 4] and [14, Remark 1], regret against the unperturbed losses and against the perturbed losses are closely related. Indeed, Corollary 9 implies the bound

$$\mathbb{E}\left[\sum_{t=1}^T F(x_t) - F(u)\right] \leq \mathbb{E}[R_T(u)] + 2D\mathbb{E}\left[\sum_{t=1}^T \|\nabla c_t(w_t)\|\right] \leq O(D\sigma\sqrt{T} + D\sqrt{GC} + 2DC).$$

4.3 Random Order Models

We apply our results from Section 3 to the Random Order model. The online ROM was introduced by [6] as a way of restricting the power of the adversary in OCO. Note that the ROM is a popular model in combinatorial online optimization, e.g., in the context of profit maximization [9]. Our results highlight that the rates in the ROM model, which is not i.i.d., are almost the same as the rates of the i.i.d. model obtained via sampling in the same set of losses with replacement.

Corollary 10. *In the single-pass ROM with convex and L -smooth losses $(f_k)_{k \in [T]}$, OFTRL (c.f. (2)) enjoys the regret bound*

$$\mathbb{E}[R_T(u)] \leq O\left(D\sigma_1 \sqrt{\log\left(e \frac{\tilde{\sigma}_1}{\sigma_1}\right) T + LD^2 + DG}\right),$$

where

$$\sigma_1^2 = \max_{x \in \mathcal{X}} \frac{1}{T} \sum_{t=1}^T \left\| \nabla f_t(x) - \frac{1}{T} \sum_{s=1}^T \nabla f_s(x) \right\|^2 \quad \text{and} \quad \tilde{\sigma}_1^2 = \frac{1}{T} \sum_{t=1}^T \max_{x \in \mathcal{X}} \left\| \nabla f_t(x) - \frac{1}{T} \sum_{s=1}^T \nabla f_s(x) \right\|^2.$$

We remark that $\sigma_1 \leq \tilde{\sigma}_1 \leq \min\{4G^2, T\sigma_1\}$. Hence, the logarithm of the ratio $\log(\tilde{\sigma}_1/\sigma_1)$ is always under control. The proof of Corollary 10 consists in controlling the adversarial variation and the cumulative variance thanks to the following lemma, proved in Appendix C.1

Lemma 11. *In the single-pass ROM, we have $\Sigma_{[1:T]}^{(2)} \leq 8G^2$ and $\bar{\sigma}_{[1:T]}^{(2)} \leq T\sigma_1^2 \log(2e^2\tilde{\sigma}_1^2/\sigma_1^2)$.*

We would like to emphasize that our results are complementary to those of [6, 31]. The focus of these works is to relax the assumption that individual losses are convex, and to only require convexity of the average loss function, leading to very different technical challenges. Inquiring if our results can also be achieved under the weaker assumptions of [31] would be an interesting direction for future work.

We also consider the multi-pass ROM. Let $P \in \mathbb{N}$ denote the number of passes. From Lemma 11 and Corollary 10 we directly obtain

$$\mathbb{E}[R_T(u)] \leq O\left(D\sigma_1 \sqrt{\log\left(e \frac{\tilde{\sigma}_1}{\sigma_1}\right) T + DG\sqrt{P} + LD^2}\right).$$

Combining Lemma [11](#) with Theorem [7](#) also gives the following corollary for strongly convex functions; see Appendix [C.2](#) for a proof.

Corollary 12. *Under the same assumption as in Theorem [7](#) the expected regret of the ROM is bounded by*

$$\mathbb{E}[R_T(u)] \leq O\left(\frac{\sigma_1^2}{\mu} \log T + \frac{G^2}{\mu} + LD^2 \kappa \log \kappa\right).$$

For multi-pass ROM with P passes, we obtain

$$\mathbb{E}[R_T(u)] \leq O\left(\frac{\sigma_1^2}{\mu} \log T + \frac{G^2}{\mu} + \frac{G^2 \log P}{n\mu} + LD^2 \kappa \log \kappa\right).$$

5 Conclusion and future work

As we showed, the exploitation of smoothness of the expected loss functions reduces the dependence of the regret bound on the maximal gradient norm to a dependence on the cumulative stochastic variance and the adversarial variation. Furthermore, we took a step towards a deeper theoretical understanding of the practically relevant intermediate scenarios. Our approach also opens several interesting new research directions. For instance, in the ROM, as mentioned in Section [4.3](#), an interesting question is whether a regret bound with dependence on σ instead of G can also be achieved with weaker assumptions as in [\[31\]](#).

Another interesting question is whether it is possible to unify the analyses and algorithms for the strongly convex and convex case. So far our analyses of these cases were intrinsically different and the choice of the algorithm requires the knowledge of strong convexity constant μ . Since this knowledge might not be available, it is of practical interest to design an adaptive method that can automatically get the best rate without manually tuning for μ . Furthermore, it would be interesting to investigate if similar or better improvements of the regret bound can be achieved with different optimistic predictions M_t . Our matching lower bounds show that there are instances where $M_t = g_{t-1}$ captures the right adversarial behavior. However, there are cases where other measures of regularity would be more relevant. One such example would be an adversary oscillating between two distributions. Interestingly, using $M_t = g_{t-2}$ would then recover optimal bounds. An interesting future research direction would be to investigate what can be achieved when using linear combinations of past gradients as optimistic predictions.

Social impact: this work is theoretical and does not give rise to any direct societal concerns.

Acknowledgments and Disclosure of Funding

Sachs, Hadiji, and Van Erven were supported by the Netherlands Organization for Scientific Research (NWO) under grant number VI.Vidi.192.095. Guzmán’s research is partially supported by the INRIA Associate Teams project, FONDECYT 1210362 grant and ANID Anillo ACT210005 grant, and National Center for Artificial Intelligence CENIA FB210017, Basal ANID. Part of this work was done while CG was at the University of Twente.

References

- [1] Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.
- [2] Idan Amir, Idan Attias, Tomer Koren, Yishay Mansour, and Roi Livni. Prediction with corrupted expert advice. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 14315–14325, 2020.
- [3] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of online learning algorithms. In *Advances in Neural Information Processing Systems 14 (NeurIPS)*, volume 14, 2002.

- [4] Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 6.1–6.20. PMLR, 25–27 Jun 2012.
- [5] Ashok Cutkosky. Anytime online-to-batch, optimism and acceleration. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 1446–1454. PMLR, 09–15 Jun 2019.
- [6] Dan Garber, Gal Korcia, and Kfir Levy. Online convex optimization in the random order model. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3387–3396. PMLR, 13–18 Jul 2020.
- [7] S. Ghadimi and G. Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM J. Optim.*, 23:2341–2368, 2013.
- [8] Saeed Ghadimi and Guanghui Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal on Optimization*, 22(4):1469–1492, 2012.
- [9] Anupam Gupta, Ruta Mehta, and Marco S. Molinaro. Maximizing profit with convex costs in the random-order model. In *ICALP*, 2018.
- [10] Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 942–953. IEEE, 2022.
- [11] Elad Hazan. Introduction to online convex optimization. *Found. Trends Optim.*, 2(3?4):157?325, aug 2016.
- [12] Elad Hazan and Satyen Kale. Extracting certainty from uncertainty: regret bounded by variation in costs. *Machine Learning*, 80:165–188, 2010.
- [13] Elad Hazan and Satyen Kale. Beyond the regret minimization barrier: an optimal algorithm for stochastic strongly-convex optimization. In *Proceedings of the 24th Annual Conference on Learning Theory*, volume 19 of *Proceedings of Machine Learning Research*, pages 421–436. PMLR, 09–11 Jun 2011.
- [14] Shinji Ito. On optimal robustness to adversarial corruption in online decision problems. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, volume 34, pages 7409–7420, 2021.
- [15] Prateek Jain, Sham M Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Accelerating stochastic gradient descent for least squares regression. In *Conference On Learning Theory*, pages 545–604. PMLR, 2018.
- [16] Pooria Joulani, András György, and Csaba Szepesvári. A modular analysis of adaptive (non-) convex optimization: Optimism, composite objectives, and variational bounds. In *International Conference on Algorithmic Learning Theory (ALT)*, pages 681–720. PMLR, 2017.
- [17] Pooria Joulani, Anant Raj, Andras Gyorgy, and Csaba Szepesvári. A simpler approach to accelerated optimization: iterative averaging meets optimism. In *International Conference on Machine Learning*, pages 4984–4993. PMLR, 2020.
- [18] Claire Kenyon. Best-fit bin-packing with random order. In *In 7th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 359–364, 1997.
- [19] Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and ℓ_1 regularization. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 525–533. PMLR, 11–13 Apr 2011.
- [20] Arkadi Nemirovski. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM J. on Optimization*, 15(1):229–251, jan 2005.

- [21] A.S. Nemirovsky and D.B. Yudin. Problem complexity and method efficiency in optimization. *SIAM Review*, 27(2):264–265, 1985.
- [22] Francesco Orabona. A modern introduction to online learning. arXiv preprint: 1912.13213, 2021.
- [23] Francesco Orabona and Dávid Pál. Scale-free online learning. *Theor. Comput. Sci.*, 716:50–69, 2018.
- [24] Boris T Polyak and Anatoli B Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- [25] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pages 993–1019. PMLR, 2013.
- [26] Alexander Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems 26 (NeurIPS)*, pages 3066–3074, 2013.
- [27] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: stochastic, constrained, and smoothed adversaries. In *Advances in Neural Information Processing Systems 24 (NeurIPS)*, pages 1764–1772, 2011.
- [28] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- [29] Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1287–1295. PMLR, 22–24 Jun 2014.
- [30] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory, Second Edition*. Society for Industrial and Applied Mathematics, USA, 2014.
- [31] Uri Sherman, Tomer Koren, and Yishay Mansour. Optimal rates for random order online optimization. In *Advances in Neural Information Processing Systems 34 (NeurIPS)*, volume 34, 2021.
- [32] Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
- [33] Tianbao Yang, Mehrdad Mahdavi, Rong Jin, and Shenghuo Zhu. Regret bounded by gradual variation for online convex optimization. *Machine Learning*, 95(2):183–223, oct 2013.
- [34] Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. In *Advances in Neural Information Processing Systems 33 (NeurIPS)*, pages 12510–12520, 2020.
- [35] Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Adaptivity and non-stationarity: Problem-dependent dynamic regret for online convex optimization. arXiv preprint arXiv:2112.14368, 2021.
- [36] Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 467–475. PMLR, 16–18 Apr 2019.
- [37] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML’03, pages 928–935. AAAI Press, 2003.