

An Efficient Algorithm for Learning with Semi-Bandit Feedback

Gergely Neu¹ and Gábor Bartók²

¹ Department of Computer Science and Information Theory
Budapest University of Technology and Economics
`gergely.neu@gmail.com`

² Department of Computer Science, ETH Zürich
`bartok@inf.ethz.ch`

Abstract. We consider the problem of online combinatorial optimization under semi-bandit feedback. The goal of the learner is to sequentially select its actions from a combinatorial decision set so as to minimize its cumulative loss. We propose a learning algorithm for this problem based on combining the Follow-the-Perturbed-Leader (FPL) prediction method with a novel loss estimation procedure called Geometric Resampling (GR). Contrary to previous solutions, the resulting algorithm can be efficiently implemented for any decision set where efficient offline combinatorial optimization is possible at all. Assuming that the elements of the decision set can be described with d -dimensional binary vectors with at most m non-zero entries, we show that the expected regret of our algorithm after T rounds is $O(m\sqrt{dT\log d})$. As a side result, we also improve the best known regret bounds for FPL in the full information setting to $O(m^{3/2}\sqrt{T\log d})$, gaining a factor of $\sqrt{d/m}$ over previous bounds for this algorithm.

Keywords: Follow-the-perturbed-leader, bandit problems, online learning, combinatorial optimization

1 Introduction

In this paper, we consider a special case of online linear optimization known as online combinatorial optimization (see Figure 1). In every time step $t = 1, 2, \dots, T$ of this sequential decision problem, the learner chooses an *action* \mathbf{V}_t from the finite action set $\mathcal{S} \subseteq \{0, 1\}^d$, where $\|\mathbf{v}\|_1 \leq m$ holds for all $\mathbf{v} \in \mathcal{S}$. At the same time, the environment fixes a loss vector $\boldsymbol{\ell}_t \in [0, 1]^d$ and the learner suffers loss $\mathbf{V}_t^\top \boldsymbol{\ell}_t$. We allow the loss vector $\boldsymbol{\ell}_t$ to depend on the previous decisions $\mathbf{V}_1, \dots, \mathbf{V}_{t-1}$ made by the learner, that is, we consider *non-oblivious* environments. The goal of the learner is to minimize the cumulative loss $\sum_{t=1}^T \mathbf{V}_t^\top \boldsymbol{\ell}_t$. Then, the performance of the learner is measured in terms of the total expected *regret*

$$R_T = \max_{\mathbf{v} \in \mathcal{S}} \mathbb{E} \left[\sum_{t=1}^T (\mathbf{V}_t - \mathbf{v})^\top \boldsymbol{\ell}_t \right] = \mathbb{E} \left[\sum_{t=1}^T \mathbf{V}_t^\top \boldsymbol{\ell}_t \right] - \min_{\mathbf{v} \in \mathcal{S}} \mathbb{E} \left[\sum_{t=1}^T \mathbf{v}^\top \boldsymbol{\ell}_t \right], \quad (1)$$

<p>Parameters: set of decision vectors $\mathcal{S} = \{\mathbf{v}(1), \mathbf{v}(2), \dots, \mathbf{v}(N)\} \subseteq \{0, 1\}^d$, number of rounds T;</p> <p>For all $t = 1, 2, \dots, T$, repeat</p> <ol style="list-style-type: none"> 1. The learner chooses a probability distribution \mathbf{p}_t over $\{1, 2, \dots, N\}$. 2. The learner draws an action I_t randomly according to \mathbf{p}_t. Consequently, the learner plays decision vector $\mathbf{V}_t = \mathbf{v}(I_t)$. 3. The environment chooses loss vector $\boldsymbol{\ell}_t$. 4. The learner suffers loss $\mathbf{V}_t^\top \boldsymbol{\ell}_t$. 5. The learner observes some feedback based on $\boldsymbol{\ell}_t$ and \mathbf{V}_t.

Fig. 1. The protocol of online combinatorial optimization.

Note that, as indicated in Figure 1, the learner chooses its actions randomly, hence the expectation.

The framework described above is general enough to accommodate a number of interesting problem instances such as path planning, ranking and matching problems, finding minimum-weight spanning trees and cut sets. Accordingly, different versions of this general learning problem have drawn considerable attention in the past few years. These versions differ in the amount of information made available to the learner after each round t . In the simplest setting, called the *full-information* setting, it is assumed that the learner gets to observe the loss vector $\boldsymbol{\ell}_t$ regardless of the choice of \mathbf{V}_t . However, this assumption does not hold for many practical applications, so it is more interesting to study the problem under *partial information*, meaning that the learner only gets some limited feedback based on its own decision. In particular, in some problems it is realistic to assume that the learner observes the vector $(V_{t,1}\ell_{t,1}, \dots, V_{t,d}\ell_{t,d})$, where $V_{t,i}$ and $\ell_{t,i}$ are the i^{th} components of the vectors \mathbf{V}_t and $\boldsymbol{\ell}_t$, respectively. This information scheme is called *semi-bandit* information. An even more challenging variant is the *full bandit* scheme where all the learner observes after time t is its own loss $\mathbf{V}_t^\top \boldsymbol{\ell}_t$.

The most well-known instance of our problem is the (adversarial) *multi-armed bandit* problem considered in the seminal paper of Auer et al. [4]: in each round of this problem, the learner has to select one of N arms and minimize regret against the best fixed arm, while only observing the losses of the chosen arm. In our framework, this setting corresponds to setting $d = N$ and $m = 1$, and assuming either full bandit or semi-bandit feedback. Among other contributions concerning this problem, Auer et al. propose an algorithm called Exp3 (Exploration and Exploitation using Exponential weights) based on constructing loss estimates $\hat{\ell}_{t,i}$ for each component of the loss vector and playing arm i with probability proportional to $\exp(-\eta \sum_{s=1}^{t-1} \hat{\ell}_{s,i})$ at time t ($\eta > 0$)³. This algorithm is known as the Exponentially Weighted Average (EWA) forecaster

³ In fact, Auer et al. mix the resulting distribution with a uniform distribution over the arms with probability $\gamma > 0$. However, this modification is not needed when one is concerned with the total expected regret.

in the full information case. Besides proving that the total expected regret of this algorithm is $O(\sqrt{NT \log N})$, Auer et al. also provide a general lower bound of $\Omega(\sqrt{NT})$ on the regret of any learning algorithm on this particular problem. This lower bound was later matched by the Implicitly Normalized Forecaster (INF) of Audibert and Bubeck [2] by using the same loss estimates in a more refined way.

The most popular example of online learning problems with actual combinatorial structure is the shortest path problem first considered by Takimoto and Warmuth [17] in the full information scheme. The same problem was considered by Györfy et al. [10], who proposed an algorithm that works with semi-bandit information. Since then, we have come a long way in understanding the “price of information” in online combinatorial optimization—see Audibert et al. [3] for a complete overview of results concerning all of the presented information schemes. The first algorithm directly targeting general online combinatorial optimization problems is due to Koolen et al. [13]: their method named Component Hedge guarantees an optimal regret of $O(m\sqrt{T \log d})$ in the full information setting. In particular, this algorithm is an instance of the more general algorithm class known as Online Stochastic Mirror Descent (OSMD) or Follow-The-Regularized-Leader (FTRL) methods. Audibert et al. [3] show that OSMD/FTRL-based methods can also be used for proving optimal regret bounds of $O(\sqrt{mdT})$ for the semi-bandit setting. Finally, Bubeck et al. [6] show that the natural extension of the EWA forecaster (coupled with an intricate exploration scheme) can be applied to obtain a $O(m^{3/2}\sqrt{dT \log d})$ upper bound on the regret when assuming full bandit feedback. This upper bound is off by a factor of $\sqrt{m \log d}$ from the lower bound proved by Audibert et al. [3]. For completeness, we note that the EWA forecaster attains a regret of $O(m^{3/2}\sqrt{T \log d})$ in the full information case and $O(m\sqrt{dT \log d})$ in the semi-bandit case.

While the results outlined above suggest that there is absolutely no work left to be done in the full information and semi-bandit schemes, we get a different picture if we restrict our attention to *computationally efficient* algorithms. First, methods based on exponential weighting of each decision vector can only be efficiently implemented for a handful of decision sets \mathcal{S} —see Koolen et al. [13] and Cesa-Bianchi and Lugosi [8] for some examples. Furthermore, as noted by Audibert et al. [3], OSMD/FTRL-type methods can be efficiently implemented by convex programming if the convex hull of the decision set can be described by a polynomial number of constraints. Details of such an efficient implementation are worked out by Suehiro et al. [16], whose algorithm runs in $O(d^6)$ time, which can still be prohibitive in practical problems. While Koolen et al. [13] list some further examples where OSMD/FTRL can be implemented efficiently, we conclude that results concerning general efficient methods for online combinatorial optimization are lacking for (semi or full) bandit information problems.

The Follow-the-Perturbed-Leader (FPL) prediction method (first proposed by Hannan [11] and later rediscovered by Kalai and Vempala [12]) method offers a computationally efficient solution for the online combinatorial optimization problem given that the *static* combinatorial optimization problem $\min_{\mathbf{v} \in \mathcal{S}} \mathbf{v}^\top \boldsymbol{\ell}$

admits computationally efficient solutions for any $\ell \in \mathbb{R}^d$. The idea underlying FPL is very simple: in every round t , the learner draws some random perturbations $\mathbf{Z}_t \in \mathbb{R}^d$ and selects the action that minimizes the perturbed total losses:

$$\mathbf{V}_t = \arg \min_{\mathbf{v} \in \mathcal{S}} \left\{ \mathbf{v}^\top \left(\sum_{s=1}^{t-1} \ell_s - \mathbf{Z}_t \right) \right\}.$$

FPL, however, is usually relatively overlooked due to many “reasons”, some of them listed below:

- The best known bound for FPL in the full information setting is $O(m\sqrt{dT})$, which is worse than the bounds for both EWA and OSMD/FTRL.
- It is commonly believed that the standard proof techniques for FPL do not apply directly against adaptive adversaries (see, e.g, the comments of Audibert et al. [3, Section 2.3] or Cesa-Bianchi and Lugosi [7, Section 4.3]). On the other hand, a direct analysis for non-oblivious adversaries is given by Poland [15] in the multi-armed bandit setting.
- Considering bandit information, no efficient FPL-style algorithm is known to achieve a regret of $O(\sqrt{T})$. Awerbuch and Kleinberg [5] and McMahan and Blum [14] proposed FPL-based algorithms for learning with full bandit feedback in shortest path problems, and proved $O(T^{2/3})$ bounds on the regret (1). Poland [15] proved bounds of $O(\sqrt{NT \log N})$ in the N -armed bandit setting, however, the proposed algorithm requires $O(T^2)$ computations per time step.

In this paper, we offer an *efficient FPL-based algorithm for regret minimization under semi-bandit feedback*. Our approach relies on a novel method for estimating components of the loss vector. The method, called *geometric resampling* (GR), is based on the idea that the reciprocal of the probability of an event can be estimated by measuring the reoccurrence time. We show that the regret of FPL coupled with GR attains a regret of $O(m\sqrt{dT \log d})$ in the semi-bandit case. To the best of our knowledge, our algorithm is the first computationally efficient learning algorithm for this learning problem. As a side result, we also improve the regret bounds of FPL in the full information setting to $O(m^{3/2}\sqrt{T \log d})$, that is, we close the gaps between the performance bounds of FPL and EWA under both full information and semi-bandit feedback.

2 Loss estimation by geometric resampling

For a gentle start, consider the problem of regret minimization in N -armed bandits where $d = N$, $m = 1$ and the learner has access to the basis vectors $\{\mathbf{e}_i\}_{i=1}^d$. In each time step, the learner specifies a distribution \mathbf{p}_t over the arms, where $p_{t,i} = \mathbb{P}[I_t = i | \mathcal{F}_{t-1}]$, where \mathcal{F}_{t-1} is the history of the learner’s observations and choices up to the end of time step $t - 1$. Most bandit algorithms rely on feeding some loss estimates to a black-box prediction algorithm. It is commonplace to

consider loss estimates of the form

$$\hat{\ell}_{t,i} = \frac{\ell_{t,i}}{p_{t,i}} \mathbb{I}\{I_t = i\}, \quad (2)$$

where $p_{t,i} = \mathbb{P}[I_t = i | \mathcal{F}_{t-1}]$, where \mathcal{F}_{t-1} is the history of observations and internal random variables used by the algorithm up to time $t - 1$. It is very easy to show that $\hat{\ell}_{t,i}$ is an unbiased estimate of the loss $\ell_{t,i}$ for all t, i such that $p_{t,i}$ is positive. For all other i and t , $\mathbb{E}[\hat{\ell}_{t,i} | \mathcal{F}_{t-1}] = 0 \leq \ell_{t,i}$.

To our knowledge, all existing bandit algorithms utilize some version of the loss estimates described above. While for many algorithms (such as the Exp3 algorithm of Auer et al. [4] and the Green algorithm of Allenberg et al. [1]), the probabilities $p_{t,i}$ are readily available and the estimates (2) can be computed efficiently, this is not necessarily the case for all algorithms. In particular, FPL is notorious for not being able to handle bandit information efficiently since the probabilities $p_{t,i}$ cannot be expressed in closed form. To overcome this difficulty, we propose a different loss estimate that can be efficiently computed *even when $p_{t,i}$ is not available for the learner*.

The estimation procedure executed after each time step t is described below.

1. The learner draws $I_t \sim \mathbf{p}_t$.
2. For $n = 1, 2, \dots$
 - (a) Let $n \leftarrow n + 1$.
 - (b) Draw $I'_t(n) \sim \mathbf{p}_t$.
 - (c) If $I'_t(n) = I_t$, break.
3. Let $K_t = n$.

Observe K_t is a geometrically distributed random variable given I_t and \mathcal{F}_{t-1} . Consequently, we have $\mathbb{E}[K_t | \mathcal{F}_{t-1}, I_t] = 1/p_{t,I_t}$. We use this property to construct the estimates

$$\hat{\ell}_{t,i} = \ell_{t,i} \mathbb{I}\{I_t = i\} K_t \quad (3)$$

for all arms i . We can easily show that the above estimate is conditionally unbiased whenever $p_{t,i} > 0$:

$$\begin{aligned} \mathbb{E}[\hat{\ell}_{t,i} | \mathcal{F}_{t-1}] &= \sum_j p_{t,j} \mathbb{E}[\hat{\ell}_{t,i} | \mathcal{F}_{t-1}, I_t = j] \\ &= p_{t,i} \mathbb{E}[\ell_{t,i} K_t | \mathcal{F}_{t-1}, I_t = i] \\ &= p_{t,i} \ell_{t,i} \mathbb{E}[K_t | \mathcal{F}_{t-1}, I_t = i] \\ &= \ell_{t,i}. \end{aligned}$$

Clearly, $\mathbb{E}[\hat{\ell}_{t,i} | \mathcal{F}_{t-1}] = 0$ still holds whenever $p_{t,i} = 0$.

The main problem with the above sampling procedure is that its worst-case running time is unbounded: while the expected number of necessary samples K_t is clearly N , the actual number of samples might be much larger. To overcome this problem, we maximize the number of samples by M and use

Algorithm 1: FPL with GR

Input: $\mathcal{S} = \{\mathbf{v}(1), \mathbf{v}(2), \dots, \mathbf{v}(N)\} \subseteq \{0, 1\}^d$, $\eta \in \mathbb{R}^+$, $M \in \mathbb{Z}^+$;
Initialization: $\widehat{\mathbf{L}}(1) = \dots = \widehat{\mathbf{L}}(d) = 0$;
for $t=1, \dots, T$ **do**
 Draw $\mathbf{Z}(1), \dots, \mathbf{Z}(d)$ independently from distribution $\text{Exp}(\eta)$;
 Choose action $I = \arg \min_{i \in \{1, 2, \dots, N\}} \{\mathbf{v}(i)^\top (\widehat{\mathbf{L}} - \mathbf{Z})\}$;
 $K(1) = \dots = K(d) = M$;
 $k = 0$; /* Counter for reoccurred indices */
 for $n=1, \dots, M-1$ **do** /* Geometric Resampling */
 Draw $\mathbf{Z}'(1), \dots, \mathbf{Z}'(d)$ independently from distribution $\text{Exp}(\eta)$;
 $I'(n) = \arg \min_{i \in \{1, 2, \dots, N\}} \{\mathbf{v}(i)^\top (\widehat{\mathbf{L}} - \mathbf{Z}')\}$;
 for $j=1, \dots, d$ **do**
 if $v(I'(n))(j) = v(I)(j) = 1$ & $K(j) = M$ **then**
 $K(j) = n$;
 $k = k + 1$;
 if $k = \|\mathbf{v}(I)\|_1$ **then break**; /* All indices reoccurred */
 end
 end
 end
 for $j=1, \dots, d$ **do** $\widehat{\mathbf{L}}(j) = \widehat{\mathbf{L}}(j) + K(j)v(I)(j)\ell(j)$; /* Update */
end

$\widetilde{K}_t = \min\{K_t, M\}$ instead of K_t in (3). While this capping obviously introduces some bias, we will show later that for appropriate values of M , this bias does not hurt the performance too much.

3 An efficient algorithm for learning with semi-bandit feedback

First, we generalize the geometric resampling method for constructing loss estimates in the semi-bandit case. To this end, let $p_{t,i} = \mathbb{P}[I_t = i | \mathcal{F}_{t-1}]$ and $q_{t,j} = \mathbb{E}[V_{t,j} | \mathcal{F}_{t-1}]$. First, the learner plays the decision vector with index $I_t \sim \mathbf{p}_t$. Then, it draws M additional indices $I'_t(1), I'_t(2), \dots, I'_t(M) \sim \mathbf{p}_t$ independently of each other and I_t . For each $j = 1, 2, \dots, d$, we define the random variables

$$K_{t,j} = \min\{1 \leq s \leq M : v_j(I'_t(s)) = 1\},$$

with the convention that $\min\{\emptyset\} = M$. We define the components of our loss estimates $\hat{\ell}_t$ as

$$\hat{\ell}_{t,j} = K_{t,j} V_{t,j} \ell_{t,j} \tag{4}$$

for all $j = 1, 2, \dots, d$. Since $V_{t,j}$ are nonzero only for coordinates for which $\ell_{t,j}$ is observed, these estimates are well-defined. It also follows that it is sufficient to

draw the additional indices $\{I'_t(s)\}$ for components with $V_{t,j} = 1$. Letting $\widehat{\mathbf{L}}_t = \sum_{s=1}^t \widehat{\ell}_s$, at time step t the algorithm draws the components of the perturbation vector \mathbf{Z}_t independently from an exponential distribution with parameter η and selects the index

$$I_t = \arg \min_{i \in \{1, 2, \dots, N\}} \left\{ \mathbf{v}(i)^\top \left(\widehat{\mathbf{L}}_{t-1} - \mathbf{Z}_t \right) \right\}.$$

As noted earlier, the distribution \mathbf{p}_t , while implicitly specified by \mathbf{Z}_t and the estimated cumulative losses $\widehat{\mathbf{L}}_t$, cannot be expressed in closed form for FPL. However, sampling the indices $I'_t(1), I'_t(2), \dots, I'_t(M)$ can be carried out by drawing additional perturbation vectors $\mathbf{Z}'_t(1), \mathbf{Z}'_t(2), \dots, \mathbf{Z}'_t(M)$ independently from the same distribution as \mathbf{Z}_t . We emphasize that the above additional indices are never actually played by the algorithm, but are only necessary for constructing the loss estimates. We also note that in general, drawing as much as M samples is usually not necessary since the sampling procedure can be terminated as soon as the values of $K_{t,i}$ are fixed for all i such that $V_{t,i} = 1$. We point the reader to Section 3.1 for a more detailed discussion of the running time of the sampling procedure.

Pseudocode for the algorithm can be found in Algorithm 1. We start analyzing our method by proving a simple lemma on the bias of the estimates.

Lemma 1. *For all $j \in \{1, 2, \dots, d\}$ and $t = 1, 2, \dots, T$ such that $q_{t,j} > 0$, the loss estimates (4) satisfy*

$$\mathbb{E} \left[\widehat{\ell}_{t,j} \mid \mathcal{F}_{t-1} \right] = \left(1 - (1 - q_{t,j})^M \right) \ell_{t,j}.$$

Proof. Fix any j, t satisfying the condition of the lemma. By elementary calculations,

$$\mathbb{E} \left[\widehat{\ell}_{t,j} \mid \mathcal{F}_{t-1} \right] = q_{t,j} \ell_{t,j} \mathbb{E} [K_{t,j} \mid \mathcal{F}_{t-1}, V_{t,j} = 1].$$

Setting $q = q_{t,j}$ for simplicity, we have

$$\begin{aligned} \mathbb{E} [K_{t,j} \mid \mathcal{F}_{t-1}, V_{t,j} = 1] &= \sum_{n=1}^{\infty} n(1-q)^{n-1}q - \sum_{n=M}^{\infty} (n-M)(1-q)^{n-1}q \\ &= \sum_{n=1}^{\infty} n(1-q)^{n-1}q - (1-q)^M \sum_{n=M}^{\infty} (n-M)(1-q)^{n-M-1}q \\ &= (1 - (1-q)^M) \sum_{n=1}^{\infty} n(1-q)^{n-1}q = \frac{1 - (1-q)^M}{q}. \end{aligned}$$

Putting the two together proves the statement. \square

The following theorem gives an upper bound on the total expected regret of the algorithm.

Theorem 1. *The total expected regret of FPL with geometric resampling satisfies*

$$R_n \leq \frac{m(\log d + 1)}{\eta} + 2\eta mdT + \frac{dT}{eM}$$

under semi-bandit information. In particular, setting $\eta = \sqrt{(\log d + 1)/(dT)}$ and $M \geq \sqrt{dT}/\left(em\sqrt{2(\log d + 1)}\right)$, the regret can be upper bounded as

$$R_n \leq 3m\sqrt{2dT(\log d + 1)}.$$

Note that the regret bound stated above holds for any non-oblivious adversary since the decision I_t only depends on the previous decisions I_{t-1}, \dots, I_1 through the loss estimates $\hat{\ell}_{t-1}, \dots, \hat{\ell}_1$. While the main ingredients of the proof presented below are rather common (we borrow several ideas from Poland [15], the proofs of Theorems 3 and 8 of Audibert et al. [3] and the proof of Corollary 4.5 in Cesa-Bianchi and Lugosi [7]), these elements are carefully combined in our proof to get the desired result.

Proof. Let $\tilde{\mathbf{Z}}$ be a perturbation vector drawn independently from the same distribution as \mathbf{Z}_1 and

$$\tilde{I}_t = \arg \min_{i \in \{1, 2, \dots, N\}} \left\{ \mathbf{v}(i)^\top \left(\hat{L}_t - \tilde{\mathbf{Z}} \right) \right\}.$$

In what follows, we will crucially use that $\tilde{\mathbf{V}}_t = \mathbf{v}(\tilde{I}_t)$ and $\mathbf{V}_{t+1} = \mathbf{v}(I_{t+1})$ are conditionally independent and identically distributed given \mathcal{F}_s for any $s \leq t$. In particular, introducing the notations

$$\begin{aligned} q_{t,k} &= \mathbb{E}[V_{t,k} | \mathcal{F}_{t-1}] & \tilde{q}_{t,k} &= \mathbb{E}[\tilde{V}_{t,k} | \mathcal{F}_t] \\ p_{t,i} &= \mathbb{P}[I_t = i | \mathcal{F}_{t-1}] & \tilde{p}_{t,i} &= \mathbb{P}[\tilde{I}_t = i | \mathcal{F}_t], \end{aligned}$$

we will exploit the above property by using $q_{t,k} = \tilde{q}_{t-1,k}$ and $p_{t,i} = \tilde{p}_{t-1,i}$ numerous times below.

First, let us address the bias of the loss estimates generated by GR. By Lemma 1, we have that $\mathbb{E}[\hat{\ell}_{t,k} | \mathcal{F}_{t-1}] \leq \ell_{t,k}$ for all k and t , and thus $\mathbb{E}[\mathbf{v}^\top \hat{\ell}_t | \mathcal{F}_{t-1}] \leq \mathbf{v}^\top \ell_t$ holds for any fixed $\mathbf{v} \in \mathcal{S}$. Furthermore, we have

$$\begin{aligned} \mathbb{E}[\tilde{\mathbf{V}}_{t-1}^\top \hat{\ell}_t | \mathcal{F}_{t-1}] &= \mathbb{E}\left[\sum_{k=1}^d \tilde{V}_{t-1,k} \hat{\ell}_{t,k} \middle| \mathcal{F}_{t-1}\right] \\ &= \sum_{k=1}^d \tilde{q}_{t-1,k} \mathbb{E}[\hat{\ell}_{t,k} | \mathcal{F}_{t-1}] \\ &= \sum_{k=1}^d \tilde{q}_{t-1,k} (1 - (1 - q_{t,k})^M) \ell_{t,k}, \end{aligned}$$

where we used the fact that $\tilde{\mathbf{V}}_{t-1}$ is independent of $\hat{\boldsymbol{\ell}}_t$ in the second line and Lemma 1 in the last line. Now using that $\tilde{q}_{t-1,k} = q_{t,k}$ for all k and t and noticing that $\mathbb{E}[\mathbf{V}_t^\top \boldsymbol{\ell}_t | \mathcal{F}_{t-1}] = \sum_{k=1}^d q_{t,k} \boldsymbol{\ell}_{t,k}$, we get that

$$\mathbb{E}[\mathbf{V}_t^\top \boldsymbol{\ell}_t | \mathcal{F}_{t-1}] \leq \mathbb{E}\left[\tilde{\mathbf{V}}_{t-1}^\top \hat{\boldsymbol{\ell}}_t | \mathcal{F}_{t-1}\right] + \sum_{i=1}^d q_{t,k} (1 - q_{t,k})^M. \quad (5)$$

To control $\sum_k q_{t,k} (1 - q_{t,k})^M$, note that $q_{t,k} (1 - q_{t,k})^M \leq q_{t,k} e^{-Mq_{t,k}}$. Since $f(q) = qe^{-Mq}$ takes its maximum at $q = 1/M$, we get

$$\sum_{k=1}^d q_{t,k} (1 - q_{t,k})^M \leq \frac{d}{eM}.$$

Using Lemma 3.1 of Cesa-Bianchi and Lugosi [7] (sometimes referred to as the “*be-the-leader*” lemma) for the sequence $(\hat{\boldsymbol{\ell}}_1 - \tilde{\mathbf{Z}}, \hat{\boldsymbol{\ell}}_2, \dots, \hat{\boldsymbol{\ell}}_T)$, we obtain

$$\sum_{t=1}^T \tilde{\mathbf{V}}_t^\top \hat{\boldsymbol{\ell}}_t - \tilde{\mathbf{V}}_1^\top \tilde{\mathbf{Z}} \leq \sum_{t=1}^T \mathbf{v}^\top \hat{\boldsymbol{\ell}}_t - \mathbf{v}^\top \tilde{\mathbf{Z}}$$

for any $\mathbf{v} \in \mathcal{S}$. Reordering and taking expectations gives

$$\mathbb{E}\left[\sum_{t=1}^T (\tilde{\mathbf{V}}_t - \mathbf{v})^\top \hat{\boldsymbol{\ell}}_t\right] \leq \mathbb{E}\left[(\tilde{\mathbf{V}}_t - \mathbf{v})^\top \tilde{\mathbf{Z}}\right] \leq \frac{m(\log d + 1)}{\eta}, \quad (6)$$

where we used $\mathbb{E}[\|\mathbf{Z}_t\|_\infty] \leq \log d + 1$. To proceed, we study the relationship between $\tilde{p}_{t,i}$ and $\tilde{p}_{t-1,i} = p_{t,i}$. To this end, we introduce the “sparse loss vector” $\hat{\boldsymbol{\ell}}'_t(i)$ with components $\hat{\ell}'_{t,k}(i) = v_k(i) \hat{\ell}_{t,k}$ and

$$\tilde{I}'_t(i) = \arg \min_{i \in \{1, 2, \dots, N\}} \left\{ \mathbf{v}(i)^\top \left(\hat{\mathbf{L}}_{t-1} + \hat{\boldsymbol{\ell}}'_t(i) - \tilde{\mathbf{Z}} \right) \right\}.$$

Using the notation $\tilde{p}'_{t,i} = \mathbb{P}\left[\tilde{I}'_t(i) = i | \mathcal{F}_t\right]$, we show in Lemma 2 (stated and proved after the proof of the theorem) that $\tilde{p}'_{t,i} \leq \tilde{p}_{t,i}$.⁴ Also, define

$$J(\mathbf{z}) = \arg \min_{j \in \{1, 2, \dots, N\}} \left\{ \mathbf{v}(j)^\top \left(\hat{\mathbf{L}}_{t-1} - \mathbf{z} \right) \right\}.$$

⁴ Note that a similar trick was used in the proof Corollary 4.5 in Cesa-Bianchi and Lugosi [7]. Also note that this trick only applies in the case of non-negative losses.

Letting $f(\mathbf{z})$ be the density of the perturbations, we have

$$\begin{aligned}
\tilde{p}_{t-1,i} &= \int_{\mathbf{z} \in [0, \infty]^d} \mathbb{I}\{J(\mathbf{z}) = i\} f(\mathbf{z}) d\mathbf{z} \\
&= e^{\eta \|\hat{\ell}'_t(i)\|_1} \int_{\mathbf{z} \in [0, \infty]^d} \mathbb{I}\{J(\mathbf{z}) = i\} f(\mathbf{z} + \hat{\ell}'_t(i)) d\mathbf{z} \\
&= e^{\eta \|\hat{\ell}'_t(i)\|_1} \int_{z_i \in [\hat{\ell}'_t(i), \infty]} \dots \int \mathbb{I}\{J(\mathbf{z} - \hat{\ell}'_t(i)) = i\} f(\mathbf{z}) d\mathbf{z} \\
&\leq e^{\eta \|\hat{\ell}'_t(i)\|_1} \int_{\mathbf{z} \in [0, \infty]^d} \mathbb{I}\{J(\mathbf{z} - \hat{\ell}'_t(i)) = i\} f(\mathbf{z}) d\mathbf{z} \\
&\leq e^{\eta \|\hat{\ell}'_t(i)\|_1} \tilde{p}'_{t,i} \leq e^{\eta \|\hat{\ell}'_t(i)\|_1} \tilde{p}_{t,i},
\end{aligned}$$

where we used $f(\mathbf{z}) = \eta \exp(-\eta \|\mathbf{z}\|_1)$ for $\mathbf{z} \in [0, \infty]^d$. Now notice that $\|\hat{\ell}'_t(i)\|_1 = \mathbf{v}(i)^\top \hat{\ell}'_t(i) = \mathbf{v}(i)^\top \hat{\ell}_t$, which yields

$$\tilde{p}_{t,i} \geq \tilde{p}_{t-1,i} e^{-\eta \mathbf{v}(i)^\top \hat{\ell}_t} \geq \tilde{p}_{t-1,i} \left(1 - \eta \mathbf{v}(i)^\top \hat{\ell}_t\right).$$

It follows that

$$\begin{aligned}
\mathbb{E} \left[\tilde{\mathbf{V}}_{t-1}^\top \hat{\ell}_t \mid \mathcal{F}_t \right] &= \sum_{i=1}^N \tilde{p}_{t-1,i} \mathbf{v}(i)^\top \hat{\ell}_t \leq \sum_{i=1}^N \tilde{p}_{t,i} \mathbf{v}(i)^\top \hat{\ell}_t + \eta \sum_{i=1}^N \tilde{p}_{t-1,i} \left(\mathbf{v}(i)^\top \hat{\ell}_t \right)^2 \\
&= \mathbb{E} \left[\tilde{\mathbf{V}}_t^\top \hat{\ell}_t \mid \mathcal{F}_t \right] + \eta \sum_{i=1}^N \tilde{p}_{t-1,i} \left(\mathbf{v}(i)^\top \hat{\ell}_t \right)^2,
\end{aligned} \tag{7}$$

where we used $\mathbb{E} \left[\tilde{\mathbf{V}}_{t-1} \mid \mathcal{F}_t \right] = \mathbb{E} \left[\tilde{\mathbf{V}}_{t-1} \mid \mathcal{F}_{t-1} \right]$ in the second equality.

Next, we will use the fact that for any j, k and t ,

$$K_{t,j} K_{t,k} \leq \frac{K_{t,j}^2 + K_{t,k}^2}{2}$$

and

$$\mathbb{E} \left[K_{t,j}^2 \mid \mathcal{F}_{t-1}, \mathbf{V}_t, \tilde{\mathbf{V}}_{t-1} \right] \leq \frac{2 - q_{t,j}}{q_{t,j}^2} \leq \frac{2}{q_{t,j}^2}$$

hold. The last term in Equation (7) can then be upper bounded as

$$\begin{aligned}
& \mathbb{E} \left[\sum_{i=1}^N \tilde{p}_{t-1,i} \left(\mathbf{v}(i)^\top \hat{\boldsymbol{\ell}}_t \right)^2 \middle| \mathcal{F}_{t-1} \right] \\
&= \mathbb{E} \left[\sum_{j=1}^d \sum_{k=1}^d \left(\tilde{V}_{t-1,j} \hat{\ell}_{t,j} \right) \left(\tilde{V}_{t-1,k} \hat{\ell}_{t,k} \right) \middle| \mathcal{F}_{t-1} \right] \\
&\leq \mathbb{E} \left[\sum_{j=1}^d \sum_{k=1}^d \frac{K_{t,j}^2 + K_{t,k}^2}{2} \left(\tilde{V}_{t-1,j} V_{t,j} \ell_{t,j} \right) \left(\tilde{V}_{t-1,k} V_{t,k} \ell_{t,k} \right) \middle| \mathcal{F}_{t-1} \right] \\
&\leq 2\mathbb{E} \left[\sum_{j=1}^d \frac{1}{q_{t,j}^2} \left(\tilde{V}_{t-1,j} V_{t,j} \ell_{t,j} \right) \sum_{k=1}^d V_{t,k} \ell_{t,k} \middle| \mathcal{F}_{t-1} \right] \\
&\leq 2m\mathbb{E} \left[\sum_{j=1}^d \ell_{t,j} \middle| \mathcal{F}_{t-1} \right] \leq 2md,
\end{aligned}$$

where in the third inequality we used that $\mathbb{E}[V_{t,j} | \mathcal{F}_{t-1}] = \mathbb{E}[\tilde{V}_{t-1,j} | \mathcal{F}_{t-1}] = q_{t,j}$. That is, we have proved

$$\mathbb{E} \left[\sum_{t=1}^T \tilde{\mathbf{V}}_{t-1}^\top \hat{\boldsymbol{\ell}}_t \right] \leq \mathbb{E} \left[\sum_{t=1}^T \tilde{\mathbf{V}}_t^\top \hat{\boldsymbol{\ell}}_t \right] + 2\eta md. \quad (8)$$

Putting Equations (5), (6) and (8) together, we obtain

$$\mathbb{E} \left[\sum_{t=1}^T (\mathbf{V}_t - \mathbf{v})^\top \boldsymbol{\ell}_t \right] \leq \frac{m(\log d + 1)}{\eta} + 2\eta mdT + \frac{dT}{eM}$$

as stated in the theorem. \square

In the next lemma, we prove that $\tilde{p}'_{t,i} \leq \tilde{p}_{t,i}$ holds for all t and i . While this statement is rather intuitive, we include its simple proof for completeness.

Lemma 2. Fix any $i \in \{1, 2, \dots, N\}$ and any vectors $\mathbf{L} \in \mathbb{R}^d$ and $\boldsymbol{\ell} \in [0, \infty)^d$. Furthermore, define the vector $\boldsymbol{\ell}'$ with components $\ell'_k = v_k(i)\ell_k$ and the perturbation vector \mathbf{Z} with independent components. Then,

$$\begin{aligned}
& \mathbb{P} \left[\mathbf{v}(i)^\top (\mathbf{L} + \boldsymbol{\ell}' - \mathbf{Z}) \leq \mathbf{v}(j)^\top (\mathbf{L} + \boldsymbol{\ell}' - \mathbf{Z}) \ (\forall j \in \{1, 2, \dots, N\}) \right] \\
& \leq \mathbb{P} \left[\mathbf{v}(i)^\top (\mathbf{L} + \boldsymbol{\ell} - \mathbf{Z}) \leq \mathbf{v}(j)^\top (\mathbf{L} + \boldsymbol{\ell} - \mathbf{Z}) \ (\forall j \in \{1, 2, \dots, N\}) \right].
\end{aligned}$$

Proof. Fix any $\forall j \in \{1, 2, \dots, N\} \setminus i$ and define the vector $\boldsymbol{\ell}'' = \boldsymbol{\ell} - \boldsymbol{\ell}'$. Define the events

$$A'_j = \{ \omega : \mathbf{v}(i)^\top (\mathbf{L} + \boldsymbol{\ell}' - \mathbf{Z}) \leq \mathbf{v}(j)^\top (\mathbf{L} + \boldsymbol{\ell}' - \mathbf{Z}) \}$$

and

$$A_j = \left\{ \omega : \mathbf{v}(i)^\top (\mathbf{L} + \boldsymbol{\ell} - \mathbf{Z}) \leq \mathbf{v}(j)^\top (\mathbf{L} + \boldsymbol{\ell} - \mathbf{Z}) \right\}.$$

We have

$$\begin{aligned} A'_j &= \left\{ \omega : (\mathbf{v}(i) - \mathbf{v}(j))^\top \mathbf{Z} \geq (\mathbf{v}(i) - \mathbf{v}(j))^\top (\mathbf{L} + \boldsymbol{\ell}') \right\} \\ &\subseteq \left\{ \omega : (\mathbf{v}(i) - \mathbf{v}(j))^\top \mathbf{Z} \geq (\mathbf{v}(i) - \mathbf{v}(j))^\top (\mathbf{L} + \boldsymbol{\ell}') - \mathbf{v}(j)^\top \boldsymbol{\ell}'' \right\} \\ &= \left\{ \omega : (\mathbf{v}(i) - \mathbf{v}(j))^\top \mathbf{Z} \geq (\mathbf{v}(i) - \mathbf{v}(j))^\top (\mathbf{L} + \boldsymbol{\ell}) \right\} = A_j, \end{aligned}$$

where we used $\mathbf{v}(i)\boldsymbol{\ell}'' = 0$ and $\mathbf{v}(j)\boldsymbol{\ell}'' \geq 0$. Now, since $A'_j \subseteq A_j$, we have $\cap_{j=1}^N A'_j \subseteq \cap_{j=1}^N A_j$, thus proving $\mathbb{P}[\cap_{j=1}^N A'_j] \leq \mathbb{P}[\cap_{j=1}^N A_j]$ as requested. \square

3.1 Running time

Let us now turn our attention to computational issues. First, we note that the efficiency of FPL-type algorithms crucially depends on the availability of an efficient oracle that solves the static combinatorial optimization problem of the form $\min_{\mathbf{v} \in \mathcal{S}} \mathbf{v}^\top \boldsymbol{\ell}$. Computing the running time of the full-information variant of FPL is straightforward: assuming that the oracle computes the solution to the static problem in $O(f(\mathcal{S}))$ time, FPL returns its prediction in $O(f(\mathcal{S}) + d)$ time (with the d overhead coming from the time necessary to generate the perturbations). Naturally, our geometric resampling scheme multiplies these computations by the number of samples taken in each round. As mentioned earlier, since we cut off the number of times we resample the decision vectors, the maximum number of additional samples per time step is M . This implies an $O((T^{3/2}d^{1/2}) \cdot (f(\mathcal{S}) + d))$ worst-case running time when using the parameter settings suggested by Theorem 1. However, the *expected* running time is much more comforting. The next simple result ensures that our algorithm requires a total of $O(f(\mathcal{S})dT + d^2T)$ processing time on expectation.

Proposition 1. *The expected number of times the algorithm draws an action up to time step T can be upper bounded by dT .*

Proof. Fix a time step t . We will denote the number of samples used by the original algorithm in round t by N_t . The expected number of samples N_t can be simply upper bounded as

$$\begin{aligned} \mathbb{E}[N_t | \mathcal{F}_{t-1}] &= \mathbb{E} \left[\max_{j: V_{t,j}=1} K_{t,j} \middle| \mathcal{F}_{t-1} \right] = \mathbb{E} \left[\max_{j=1,2,\dots,d} V_{t,j} K_{t,j} \middle| \mathcal{F}_{t-1} \right] \\ &\leq \mathbb{E} \left[\sum_{j=1}^d V_{t,j} K_{t,j} \middle| \mathcal{F}_{t-1} \right]. \end{aligned}$$

By Lemma 1, we have $\mathbb{E}[K_{t,j} | \mathcal{F}_{t-1}, V_{t,j}] \leq 1/\mathbb{E}[V_{t,j} | \mathcal{F}_{t-1}]$, so we obtain $\mathbb{E}[N_t] \leq d$ and the statement of the lemma follows. \square

4 Improved bounds for learning with full information

Our proof technique also enables us to tighten the guarantees for FPL in the full information setting. In particular, we consider the algorithm choosing the index

$$I_t = \arg \min_{i \in \{1, 2, \dots, N\}} \{ \mathbf{v}(i)^\top (\mathbf{L}_{t-1} - \mathbf{Z}_t) \},$$

where $\mathbf{L}_t = \sum_{s=1}^t \boldsymbol{\ell}_s$ and the components of \mathbf{Z}_t are drawn independently from an exponential distribution with parameter η . We state our improved regret bounds concerning this algorithm in the following theorem.

Theorem 2. *Let $C_T = \sum_{t=1}^T \mathbb{E} [\mathbf{V}_t^\top \boldsymbol{\ell}_t]$. Then the total expected regret of FPL satisfies*

$$R_n \leq \frac{m(\log d + 1)}{\eta} + \eta m C_T$$

under full information. In particular, setting $\eta = \sqrt{(\log d + 1)/(mT)}$, the regret can be upper bounded as

$$R_n \leq 2m^{3/2} \sqrt{T(\log d + 1)}.$$

Note that the above bound can be further tightened if some upper bound $C_T^* \geq C_T$ is available a priori. Once again, these regret bounds hold for any non-oblivious adversary since the decision I_t depends on the previous decisions I_{t-1}, \dots, I_1 only through the loss vectors $\boldsymbol{\ell}_{t-1}, \dots, \boldsymbol{\ell}_1$. The bound is an improvement upon the previous known bound for FPL by Kalai and Vempala [12], which is of $O(m\sqrt{dT})$.

Proof. The statement follows from a simplification of the proof of Theorem 1 when using $\hat{\boldsymbol{\ell}}_t = \boldsymbol{\ell}_t$. First, identically to Equation (6), we have

$$\mathbb{E} \left[\sum_{t=1}^T (\tilde{\mathbf{V}}_t - \mathbf{v})^\top \boldsymbol{\ell}_t \right] \leq \mathbb{E} \left[(\tilde{\mathbf{V}}_t - \mathbf{v})^\top \tilde{\mathbf{Z}} \right] \leq \frac{m(\log d + 1)}{\eta}.$$

Further, it is easy to see that the conditions of Lemma 2 are satisfied and, similarly to Equation (7), we also have

$$\begin{aligned} \mathbb{E} \left[\tilde{\mathbf{V}}_{t-1}^\top \boldsymbol{\ell}_t \right] &\leq \mathbb{E} \left[\tilde{\mathbf{V}}_t^\top \boldsymbol{\ell}_t \right] + \eta \sum_{i=1}^N \tilde{p}_{t-1,i} (\mathbf{v}(i)^\top \boldsymbol{\ell}_t)^2 \\ &\leq \mathbb{E} \left[\tilde{\mathbf{V}}_t^\top \boldsymbol{\ell}_t \right] + \eta m \sum_{i=1}^N \tilde{p}_{t-1,i} \mathbf{v}(i)^\top \boldsymbol{\ell}_t. \end{aligned}$$

Using that \mathbf{V}_t and $\tilde{\mathbf{V}}_{t-1}$ have the same distribution, we obtain the statement of the theorem. \square

5 Conclusions and open problems

In this paper, we have described the first general efficient algorithm for online combinatorial optimization under semi-bandit feedback. We have proved that the regret of our algorithm is $O(m\sqrt{dT\log d})$ in this setting, and have also shown that FPL can achieve $O(m^{3/2}\sqrt{T\log d})$ in the full information case when tuned properly. While these bounds are off by a factor of $\sqrt{m\log d}$ and \sqrt{m} from the respective minimax results, they exactly match the best known regret bounds for the well-studied Exponentially Weighted Forecaster (EWA). Whether the gaps mentioned above can be closed for FPL-style algorithms (e.g., by using more intricate perturbation schemes) remains an important open question. Nevertheless, we regard our contribution as a significant step towards understanding the inherent trade-offs between computational efficiency and performance guarantees in online combinatorial optimization and, more generally, in online linear optimization.

The efficiency of our method rests on a novel loss estimation method called geometric resampling (GR). Obviously, this estimation method is not specific to the proposed learning algorithm. While GR has no immediate benefits for OSMD/FTRL-type algorithms where the probabilities $q_{t,k}$ are readily available, it is possible to think about problem instances where EWA can be efficiently implemented while the values of $q_{t,k}$ are difficult to compute.

The most important open problem left is the case of efficient online linear optimization with full bandit feedback. Learning algorithms for this problem usually require that the pseudoinverse of the covariance matrix $P_t = \mathbb{E}[\mathbf{V}_t\mathbf{V}_t^\top | \mathcal{F}_{t-1}]$ is readily available for the learner at each time step (see, e.g., McMahan and Blum [14], Dani et al. [9], Cesa-Bianchi and Lugosi [8], Bubeck et al. [6]). While for most problems, this inverse matrix cannot be computed efficiently, it can be efficiently approximated by geometric resampling when P_t is positive definite as the limit of the matrix geometric series $\sum_{n=1}^{\infty} (I - P_t)^n$. While this knowledge should be enough to construct an efficient FPL-based method for online combinatorial optimization under full bandit feedback, we have to note that the analysis presented in this paper does not carry through directly in this case: as usual loss estimates might take negative values in the full bandit setting, proving a bound similar to Equation (7) cannot be performed in the presented manner.

Acknowledgments This work was supported in part by the Hungarian Scientific Research Fund and the Hungarian National Office for Research and Technology (KTIA-OTKA CNK 77782) and by DARPA grant MSEE FA8650-11-1-7156.

References

- [1] Allenberg, C., Auer, P., Györfi, L., and Ottucsák, Gy. (2006). Hannan consistency in on-line learning in case of unbounded losses under partial monitoring. In *ALT*, pages 229–243.
- [2] Audibert, J.-Y. and Bubeck, S. (2010). Regret bounds and minimax policies under partial monitoring. *Journal of Machine Learning Research*, 11:2635–2686.
- [3] Audibert, J. Y., Bubeck, S., and Lugosi, G. (2013). Regret in online combinatorial optimization. *To appear in Mathematics of Operations Research*.
- [4] Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2002). The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77.
- [5] Awerbuch, B. and Kleinberg, R. D. (2004). Adaptive routing with end-to-end feedback: distributed learning and geometric approaches. In *Proceedings of the 36th ACM Symposium on Theory of Computing*, pages 45–53.
- [6] Bubeck, S., Cesa-Bianchi, N., and Kakade, S. M. (2012). Towards minimax policies for online linear optimization with bandit feedback. *Journal of Machine Learning Research - Proceedings Track*, 23:41.1–41.14.
- [7] Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA.
- [8] Cesa-Bianchi, N. and Lugosi, G. (2012). Combinatorial bandits. *Journal of Computer and System Sciences*, 78:1404–1422.
- [9] Dani, V., Hayes, T., and Kakade, S. (2008). The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems (NIPS)*, volume 20, pages 345–352.
- [10] György, A., Linder, T., Lugosi, G., and Ottucsák, Gy. (2007). The on-line shortest path problem under partial monitoring. *Journal of Machine Learning Research*, 8:2369–2403.
- [11] Hannan, J. (1957). Approximation to Bayes risk in repeated play. *Contributions to the theory of games*, 3:97–139.
- [12] Kalai, A. and Vempala, S. (2005). Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71:291–307.
- [13] Koolen, W., Warmuth, M., and Kivinen, J. (2010). Hedging structured concepts. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 93–105.
- [14] McMahan, H. B. and Blum, A. (2004). Online geometric optimization in the bandit setting against an adaptive adversary. In *Proceedings of the Eighteenth Conference on Computational Learning Theory*, pages 109–123.
- [15] Poland, J. (2005). FPL analysis for adaptive bandits. In *In 3rd Symposium on Stochastic Algorithms, Foundations and Applications (SAGA'05)*, pages 58–69.
- [16] Suehiro, D., Hatano, K., Kijima, S., Takimoto, E., and Nagano, K. (2012). Online prediction under submodular constraints. In *Algorithmic Learning Theory*, volume 7568 of *Lecture Notes in Computer Science*, pages 260–274. Springer Berlin Heidelberg.
- [17] Takimoto, E. and Warmuth, M. (2003). Paths kernels and multiplicative updates. *Journal of Machine Learning Research*, 4:773–818.