

Prediction by random-walk perturbation

Luc Devroye

*School of Computer Science
McGill University*

LUCDEVROYE@GMAIL.COM

Gábor Lugosi

*ICREA and Department of Economics
Universitat Pompeu Fabra*

GABOR.LUGOSI@GMAIL.COM

Gergely Neu

*MTA SZTAKI Institute for Computer Science and Control
and*

*Department of Computer Science and Information Theory
Budapest University of Technology and Economics*

GERGELY.NEU@GMAIL.COM

Abstract

We propose a version of the follow-the-perturbed-leader online prediction algorithm in which the cumulative losses are perturbed by independent symmetric random walks. The forecaster is shown to achieve an expected regret of the optimal order $O(\sqrt{n \log N})$ where n is the time horizon and N is the number of experts. More importantly, it is shown that the forecaster changes its prediction at most $O(\sqrt{n \log N})$ times, in expectation. We also extend the analysis to online combinatorial optimization and show that even in this more general setting, the forecaster rarely switches between experts while having a regret of near-optimal order.

Keywords: Online learning, Online combinatorial optimization, Follow the Perturbed Leader, Random walk

1. Preliminaries

In this paper we study the problem of online prediction with expert advice, see [Cesa-Bianchi and Lugosi \(2006\)](#). The problem may be described as a repeated game between a *forecaster* and an adversary—the *environment*. At each time instant $t = 1, \dots, n$, the forecaster chooses one of the N available actions (often called *experts*) and suffers a loss $\ell_{i,t} \in [0, 1]$ corresponding to the chosen action i . We consider the so-called *oblivious adversary* model in which the environment selects all losses before the prediction game starts and reveals the losses $\ell_{i,t}$ at time t after the forecaster has made its prediction. The losses are deterministic but the forecaster may randomize: at time t , the forecaster chooses a probability distribution \mathbf{p}_t over the set of N actions and draws a random action I_t according to the distribution \mathbf{p}_t . The prediction protocol is described in Figure 1.

The usual goal for the standard prediction problem is to devise an algorithm such that the cumulative loss $\widehat{L}_n = \sum_{t=1}^n \ell_{I_t,t}$ is as small as possible, in expectation and/or with high

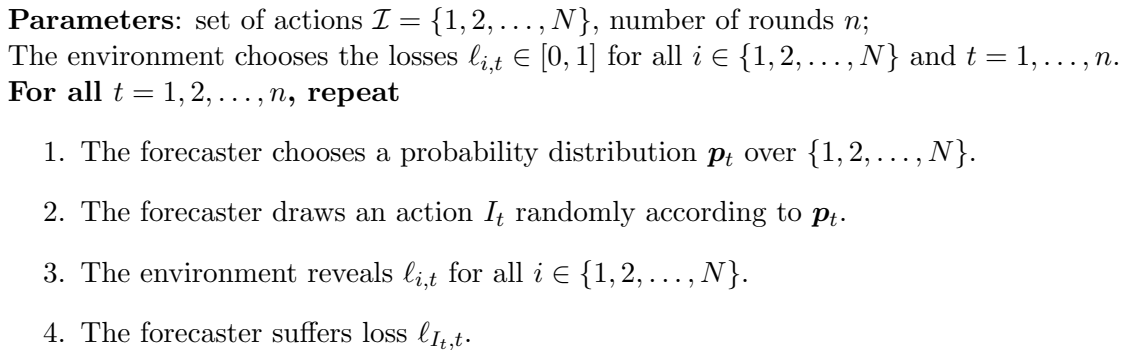


Figure 1: Prediction with expert advice.

probability (where probability is with respect to the forecaster's randomization). Since we do not make any assumption on how the environment generates the losses $\ell_{i,t}$, we cannot hope to minimize the above cumulative loss. Instead, a meaningful goal is to minimize the performance gap between our algorithm and the strategy that selects the best action chosen in hindsight. This performance gap is called the *regret* and is defined formally as

$$R_n = \max_{i \in \{1, 2, \dots, N\}} \sum_{t=1}^n (\ell_{I_t,t} - \ell_{i,t}) = \widehat{L}_n - L_n^*,$$

where we have also introduced the notation $L_n^* = \min_{i \in \{1, 2, \dots, N\}} \sum_{t=1}^n \ell_{i,t}$. Minimizing the regret defined above is a well-studied problem. It is known that no matter what algorithm the forecaster uses,

$$\liminf_{n, N \rightarrow \infty} \sup \frac{\mathbb{E}R_n}{\sqrt{(n/2) \ln N}} \geq 1$$

where the supremum is taken with respect to all possible loss assignments with losses in $[0, 1]$ (see, e.g., [Cesa-Bianchi and Lugosi \(2006\)](#)). On the other hand, several prediction algorithms are known whose expected regret is of optimal order $O(\sqrt{n \log N})$ and many of them achieve a regret of this order with high probability. Perhaps the most popular one is the exponentially weighted average forecaster (a variant of weighted majority algorithm of [Littlestone and Warmuth \(1994\)](#), and aggregating strategies of [Vovk \(1990\)](#), also known as *Hedge* by [Freund and Schapire \(1997\)](#)). The exponentially weighted average forecaster assigns probabilities to the actions that are inversely proportional to an exponential function of the loss accumulated by each action up to time t .

Another popular forecaster is the *follow the perturbed leader* (FPL) algorithm of [Hannan \(1957\)](#). [Kalai and Vempala \(2003\)](#) showed that Hannan's forecaster, when appropriately modified, indeed achieves an expected regret of optimal order. At time t , the FPL forecaster adds a random perturbation $Z_{i,t}$ to the cumulative loss $L_{i,t-1} = \sum_{s=1}^{t-1} \ell_{i,s}$ of each action and chooses an action that minimizes the sum $L_{i,t-1} + Z_{i,t}$. If the vector of random variables $\mathbf{Z}_t = (Z_{1,t}, \dots, Z_{N,t})$ have joint density $(\eta/2)^N e^{-\eta \|\mathbf{z}\|_1}$ for $\eta \sim \sqrt{\log N/n}$, then the expected regret of the forecaster is of order $O(\sqrt{n \log N})$ ([Kalai and Vempala \(2005\)](#), see also [Cesa-Bianchi and Lugosi \(2006\)](#), [Hutter and Poland \(2004\)](#), [Poland \(2005\)](#)). This is true whether $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ are independent or not. If they are independent, then one may show that the

regret is concentrated around its expectation. Another interesting choice is when $\mathbf{Z}_1 = \dots = \mathbf{Z}_n$, that is, the same perturbation is used over time. Even though this forecaster has an expected regret of optimal order, the regret is much less concentrated and may fail with reasonably high probability.

Small regret is not the only desirable feature of an online forecasting algorithm. In many applications, one would like to define forecasters that do not change their prediction too often. Examples of such problems include the online buffering problem described by Geulen et al. (2010) and the online lossy source coding problem of György and Neu (2011). A more abstract problem where the number of abrupt switches in the behavior is costly is the problem of online learning in Markovian decision processes, as described by Even-Dar et al. (2009) and Neu et al. (2010).

To be precise, define the number of action switches up to time n by

$$C_n = |\{1 < t \leq n : I_{t+1} \neq I_t\}| .$$

In particular, we are interested in defining randomized forecasters that achieve a regret R_n of the order $O(\sqrt{n \log N})$ while keeping the number of action switches C_n as small as possible. However, the usual forecasters with small regret—such as the exponentially weighted average forecaster or the FPL forecaster with i.i.d. perturbations—may switch actions a large number of times, typically $\Theta(n)$. Therefore, the design of special forecasters with small regret and small number of action switches is called for.

The first paper to explicitly attack this problem is by Geulen et al. (2010), who propose a variant of the exponentially weighted average forecaster called the “Shrinking Dartboard” algorithm and prove that it provides an expected regret of $O(\sqrt{n \log N})$, while guaranteeing that the expected number of switches is at most $O(\sqrt{n \log N})$. A less conscious attempt to solve the problem is due to Kalai and Vempala (2005); they show that the simplified version of the FPL algorithm with identical perturbations (as described above) guarantees an $O(\sqrt{n \log N})$ bound on both the expected regret and the expected number of switches. In this paper, we propose a method based on FPL in which perturbations are defined by independent symmetric random walks. We show that this, intuitively appealing, forecaster has similar regret and switch-number guarantees as Shrinking Dartboard and FPL with identical perturbations. A further important advantage of the new forecaster is that it may be used simply in the more general problem of *online combinatorial*—or, more generally, *linear—optimization*. We postpone the definitions and the statement of the results to Section 4 below.

2. The algorithm

To address the problem described in the previous section, we propose a variant of the Follow the Perturbed Leader (FPL) algorithm. The proposed forecaster perturbs the loss of each action at every time instant by a symmetric coin flip and chooses an action with minimal cumulative perturbed loss. More precisely, the algorithm draws independent random variables $X_{i,t}$ that take values $\pm 1/2$ with equal probabilities and $X_{i,t}$ is added to each loss $\ell_{i,t-1}$. At time t action i is chosen that minimizes $\sum_{s=1}^t (\ell_{i,t-1} + X_{i,t})$ (where we define $\ell_{i,0} = 0$).

Equivalently, the forecaster may be thought of as an FPL algorithm in which the cumulative losses $L_{i,t-1}$ are perturbed by $Z_{i,t} = \sum_{s=1}^t X_{i,s}$. Since for each fixed i , $Z_{i,1}, Z_{i,2}, \dots$ is a

Algorithm 1 Prediction by random-walk perturbation.

Initialization: set $L_{i,0} = 0$ and $Z_{i,0} = 0$ for all $i = 1, 2, \dots, N$.

For all $t = 1, 2, \dots, n$, **repeat**

1. Draw $X_{i,t}$ for all $i = 1, 2, \dots, N$ such that

$$X_{i,t} = \begin{cases} \frac{1}{2} & \text{with probability } \frac{1}{2} \\ -\frac{1}{2} & \text{with probability } \frac{1}{2}. \end{cases}$$

2. Let $Z_{i,t} = Z_{i,t-1} + X_{i,t}$ for all $i = 1, 2, \dots, N$.

3. Choose action

$$I_t = \arg \min_i (L_{i,t-1} + Z_{i,t}),$$

where ties are broken in favor of I_{t-1} .

4. Observe losses $\ell_{i,t}$ for all $i = 1, 2, \dots, N$, suffer loss $\ell_{I_t,t}$.

5. Set $L_{i,t} = L_{i,t-1} + \ell_{i,t}$ for all $i = 1, 2, \dots, N$.
-

symmetric random walk, cumulative losses of the N actions are perturbed by N independent symmetric random walks. This is the way the algorithm is presented in Algorithm 1.

A simple variation is when one replaces random coin flips by independent standard normal random variables. Both have similar performance guarantees and we choose $\pm(1/2)$ -valued perturbations for mathematical convenience. In Section 4 we switch to normally distributed perturbations—again driven by mathematical simplicity. In practice both versions are expected to have a similar behavior.

Conceptually, the difference between standard FPL and the proposed version is the way the perturbations are generated: while common versions of FPL use perturbations that are generated in an i.i.d. fashion, the perturbations of the algorithm proposed here are dependent. This will enable us to control the number of action switches during the learning process. Note that the standard deviation of these perturbations is still of order \sqrt{t} just like for the standard FPL forecaster with optimal parameter settings.

To obtain intuition why this approach will solve our problem, first consider a problem with $N = 2$ actions and an environment that generates equal losses, say $\ell_{i,t} = 0$ for all i and t . When using i.i.d. perturbations, FPL switches actions with probability $1/2$ in each round, thus yielding $C_t = t/2 + O(\sqrt{t})$ with overwhelming probability. The same holds for the exponentially weighted average forecaster. On the other hand, when using the random-walk perturbations described above, we only switch between the actions when the leading random walk is changed, that is, when the difference of the two random walks—which is also a symmetric random walk—hits zero. It is a well known that the number of occurrences of this event up to time t is $O_p(\sqrt{t})$, see, [Feller \(1968\)](#). As we show below, this is the worst case for the number of switches.

A similar variant of the FPL forecaster was recently derived by [Rakhlin et al. \(2012\)](#), who use perturbations of the form $Z_{i,t} = \sum_{s=t+1}^n X_{i,s}$, where $X_{i,s}$ are i.i.d. random variables with

an arbitrary symmetric distribution. [Rakhlin et al.](#) exploit the fact that these perturbations can serve as a relaxation of the Rademacher complexity of the prediction game, and prove an $O(\sqrt{n \log N})$ bound on the expected regret of the resulting algorithm. While this approach cannot be used for analyzing our FPL-variant, our analysis can be directly applied to provide both regret and switch-number guarantees for their method. Additionally, note that our algorithm does not need to use prior knowledge of the number of rounds.

3. Performance bounds

The next theorem summarizes our performance bounds for the proposed forecaster.

Theorem 1 *The expected regret and expected number of switches of actions of the forecaster of [Algorithm 1](#) satisfy, for all possible loss sequences (under the oblivious-adversary model),*

$$\mathbb{E}R_n \leq 2\mathbb{E}C_n \leq 8\sqrt{2n \log N} + 16 \log n + 16 .$$

Even though we only prove bounds for the expected regret and the expected number of switches, it is of great interest to understand upper tail probabilities. However, this is a highly nontrivial problem. One may get an intuition by considering the case when $N = 2$ and all losses are equal to zero. In this case the algorithm switches actions whenever a symmetric random walk returns to zero. This distribution is well understood and the probability that this occurs more than $x\sqrt{n}$ times during the first n steps is roughly $2\mathbb{P}\{N > 2x\} \leq 2e^{-2x^2}$ where N is a standard normal random variable (see [Feller, 1968](#), Section III.4). Thus, in this case we see that the number of switches is bounded by $O\left(\sqrt{n \log(1/\delta)}\right)$, with probability at least $1 - \delta$. However, proving analogous bounds for the general case remains a challenge.

To prove the theorem, we first show that the regret can be bounded in terms of the number of action switches. Then we turn to analyzing the expected number of action switches.

3.1. Regret and the number of switches

The next simple lemma shows that the regret of the forecaster may be bounded in terms of the number of times the forecaster switches actions.

Lemma 2 *Fix any $i \in \{1, 2, \dots, N\}$. Then*

$$\widehat{L}_n - L_{i,n} \leq 2C_n + Z_{i,n+1} - \sum_{t=1}^{n+1} X_{I_{t-1},t} .$$

Proof We apply Lemma 3.1 of [Cesa-Bianchi and Lugosi \(2006\)](#) (sometimes referred to as the “be-the-leader” lemma) for the sequence $(\ell_{i,t-1} + X_{i,t})_{t=1}^{\infty}$ with $\ell_{j,0} = 0$ for all $j \in \{1, 2, \dots, N\}$, obtaining

$$\begin{aligned} \sum_{t=1}^{n+1} (\ell_{I_t,t-1} + X_{I_t,t}) &\leq \sum_{t=1}^{n+1} (\ell_{i,t-1} + X_{i,t}) \\ &= L_{i,n} + Z_{i,n+1} . \end{aligned}$$

Reordering terms, we get

$$\sum_{t=1}^n \ell_{I_t,t} \leq L_{i,n} + \sum_{t=1}^{n+1} (\ell_{I_{t-1},t-1} - \ell_{I_t,t-1}) + Z_{i,n} - \sum_{t=1}^{n+1} X_{I_t,t} . \quad (1)$$

The last term can be rewritten as

$$-\sum_{t=1}^{n+1} X_{I_t,t} = -\sum_{t=1}^{n+1} X_{I_{t-1},t} + \sum_{t=1}^{n+1} (X_{I_{t-1},t} - X_{I_t,t}) .$$

Now notice that $X_{I_{t-1},t} - X_{I_t,t}$ and $\ell_{I_{t-1},t-1} - \ell_{I_t,t-1}$ are both zero when $I_t = I_{t-1}$ and are upper bounded by 1 otherwise. That is, we get that

$$\sum_{t=1}^{n+1} (\ell_{I_{t-1},t-1} - \ell_{I_t,t-1}) + \sum_{t=1}^{n+1} (X_{I_{t-1},t} - X_{I_t,t}) \leq 2 \sum_{t=1}^{n+1} \mathbb{I}\{I_{t-1} \neq I_t\} = 2C_n .$$

Putting everything together gives the statement of the lemma. \blacksquare

3.2. Bounding the number of switches

Next we analyze the number of switches C_n . In particular, we upper bound the marginal probability $\mathbb{P}[I_{t+1} \neq I_t]$ for each $t \geq 1$. We define the *lead pack* A_t as the set of actions that, at time t , have a positive probability of taking the lead at time $t+1$:

$$A_t = \left\{ i \in \{1, 2, \dots, N\} : L_{i,t-1} + Z_{i,t} < \min_j (L_{j,t-1} + Z_{j,t}) + 2 \right\} ,$$

where we assumed that ties are broken in favor of I_{t-1} . We bound the probability of lead change as

$$\mathbb{P}[I_t \neq I_{t+1}] \leq \frac{1}{2} \mathbb{P}[|A_t| > 1] .$$

The key to the proof of the theorem is the following lemma that gives an upper bound for the probability that the lead pack contains more than one action. It implies, in particular, that

$$\mathbb{E}[C_n] \leq 4\sqrt{2n \log N} + 4 \log n + 4 ,$$

which is what we need to prove the expected-value bounds of Theorem 1.

Lemma 3

$$\mathbb{P}[|A_t| > 1] \leq 4\sqrt{2\frac{\log N}{t}} + \frac{8}{t} .$$

Proof Define $p_t(k) = \mathbb{P}[Z_{1,t} = \frac{k}{2}]$ for all $k = -t, \dots, t$ and we let S_t denote the set of leaders at time t :

$$S_t = \left\{ j \in \{1, 2, \dots, N\} : L_{j,t-1} + Z_{j,t} = \min_i \{L_{i,t-1} + Z_{i,t}\} \right\} .$$

The forecaster picks $I_t \in S_t$ arbitrarily when $I_{t-1} \notin S_t$, otherwise it stays with $I_t = I_{t-1}$. Let us start with analyzing $\mathbb{P}[|A_t| = 1]$:

$$\begin{aligned} \mathbb{P}[|A_t| = 1] &= \sum_{k=-t}^t \sum_{j=1}^N p_t(k) \mathbb{P} \left[\min_{i \neq j} \{L_{i,t-1} + Z_{i,t}\} \geq L_{j,t-1} + \frac{k}{2} + 2 \right] \\ &\geq \sum_{k=-t}^{t-4} \sum_{j=1}^N p_t(k+4) \mathbb{P} \left[\min_{i \neq j} \{L_{i,t-1} + Z_{i,t}\} \geq L_{j,t-1} + \frac{k+4}{2} \right] \frac{p_t(k)}{p_t(k+4)} \\ &= \sum_{k=-t+4}^t \sum_{j=1}^N p_t(k) \mathbb{P} \left[\min_{i \neq j} \{L_{i,t-1} + Z_{i,t}\} \geq L_{j,t-1} + \frac{k}{2} \right] \frac{p_t(k-4)}{p_t(k)}. \end{aligned}$$

Before proceeding, we need to make two observations. First of all,

$$\begin{aligned} \sum_{j=1}^N p_t(k) \mathbb{P} \left[\min_{i \neq j} \{L_{i,t-1} + Z_{i,t}\} \geq L_{j,t-1} + \frac{k}{2} \right] &\geq \mathbb{P} \left[\exists j \in S_t : Z_{j,t} = \frac{k}{2} \right] \\ &\geq \mathbb{P} \left[\min_{j \in S_t} Z_{j,t} = \frac{k}{2} \right], \end{aligned}$$

where the first inequality follows from the union bound and the second from the fact that the latter event implies the former. Also notice that $\frac{t+Z_{1,t}}{2}$ is binomially distributed with parameters t and $1/2$ and therefore $p_t(k) = \binom{t}{\frac{t+k}{2}} \frac{1}{2^t}$. Hence

$$\begin{aligned} \frac{p_t(k-4)}{p_t(k)} &= \frac{\binom{t+k}{2}! \binom{t-k}{2}!}{\binom{t+k}{2}-2)! \binom{t-k}{2}+2)!} \\ &= 1 + \frac{4(t+1)(k-2)}{(t-k+2)(t-k+4)}. \end{aligned}$$

It can be easily verified that

$$\frac{4(t+1)(k-2)}{(t-k+2)(t-k+4)} \geq \frac{4(t+1)(k-2)}{(t+2)(t+4)}$$

holds for all $k \in [-t, t]$. Using our first observation, we get

$$\begin{aligned} \mathbb{P}[|A_t| = 1] &\geq \sum_j \sum_{k=-t+4}^t p_t(k) \mathbb{P} \left[\min_{i \neq j} \{L_{i,t-1} + Z_{i,t}\} \geq L_{j,t-1} + \frac{k}{2} \right] \frac{p_t(k-4)}{p_t(k)} \\ &\geq \sum_{k=-t+4}^t \mathbb{P} \left[\min_{j \in S_t} Z_{j,t} = \frac{k}{2} \right] \frac{p_t(k-4)}{p_t(k)}. \end{aligned}$$

Along with our second observation, this implies

$$\begin{aligned}
\mathbb{P}[|A_t| > 1] &\leq 1 - \sum_{k=-t+4}^t \mathbb{P}\left[\min_{j \in S_t} Z_{j,t} = \frac{k}{2}\right] \frac{p_t(k-4)}{p_t(k)} \\
&\leq 1 - \sum_{k=-t+4}^t \mathbb{P}\left[\min_{j \in S_t} Z_{j,t} = \frac{k}{2}\right] \left(1 + \frac{4(t+1)(k-2)}{(t+2)(t+4)}\right) \\
&\leq \sum_{k=-t}^t \mathbb{P}\left[\min_{j \in S_t} Z_{j,t} = \frac{k}{2}\right] \left(\frac{4(2-k)(t+1)}{(t+2)(t+4)}\right) \\
&= \frac{8(t+1)}{(t+2)(t+4)} - 8 \frac{t+1}{(t+2)(t+4)} \mathbb{E}\left[\min_{j \in S_t} Z_{j,t}\right] \\
&\leq \frac{8}{t} + \frac{8}{t} \mathbb{E}\left[\max_{j \in \{1,2,\dots,N\}} Z_{j,t}\right].
\end{aligned}$$

Now using $\mathbb{E}[\max_j Z_{j,t}] \leq \sqrt{\frac{t \log N}{2}}$ implies

$$\mathbb{P}[|A_t| > 1] \leq 4\sqrt{\frac{2 \log N}{t}} + \frac{8}{t}$$

as desired. ■

4. Online combinatorial optimization

In this section we study the problem of online linear optimization (see, among others, [Gentile and Warmuth \(1998\)](#), [Kivinen and Warmuth \(2001\)](#), [Grove et al. \(2001\)](#), [Takimoto and Warmuth \(2003\)](#), [Kalai and Vempala \(2005\)](#), [Warmuth and Kuzmin \(2008\)](#), [Helmbold and Warmuth \(2009\)](#), [Hazan et al. \(2010\)](#), [Koolen et al. \(2010\)](#), [Cesa-Bianchi and Lugosi \(2012\)](#), [Audibert et al. \(2011\)](#)). This is a similar prediction problem as the one described in the introduction but here each action i is represented by a vector $\mathbf{v}_i \in \mathbb{R}^d$. The loss corresponding to action i at time t equals $\mathbf{v}_i^\top \boldsymbol{\ell}_t$ where $\boldsymbol{\ell}_t \in [0, 1]^d$ is the so-called *loss vector*. Thus, given a set of actions $\mathcal{S} = \{\mathbf{v}_i : i = 1, 2, \dots, N\} \subseteq \mathbb{R}^d$, at every time instant t , the forecaster chooses, in a possibly randomized way, a vector $\mathbf{V}_t \in \mathcal{S}$ and suffers loss $\mathbf{V}_t^\top \boldsymbol{\ell}_t$. We denote by $\hat{L}_n = \sum_{t=1}^n \mathbf{V}_t^\top \boldsymbol{\ell}_t$ the cumulative loss of the forecaster and the regret becomes

$$\hat{L}_n - \min_{\mathbf{v} \in \mathcal{S}} \mathbf{v}^\top \mathbf{L}_n$$

where $\mathbf{L}_t = \sum_{s=1}^t \boldsymbol{\ell}_s$ is the cumulative loss vector. While the results of the previous section still hold when treating each $\mathbf{v}_i \in \mathcal{S}$ as a separate action, one may gain important computational advantage by taking the structure of the action set into account. In particular, as [Kalai and Vempala \(2005\)](#) emphasize, FPL-type forecasters may often be computed efficiently. In this section we propose such a forecaster which adds independent random-walk perturbations to the *individual components* of the loss vector. To gain simplicity in the presentation, we restrict our attention to the case of *online combinatorial optimization* in

which $\mathcal{S} \subset \{0, 1\}^d$, that is, each action is represented a binary vector. This special case arguably contains most important applications such as the *online shortest path* problem. In this example, a fixed directed acyclic graph of d edges is given with two distinguished vertices u and w . The forecaster, at every time instant t , chooses a directed path from u to w . Such a path is represented by its binary incidence vector $\mathbf{v} \in \{0, 1\}^d$. The components of the loss vector $\boldsymbol{\ell}_t \in [0, 1]^d$ represent losses assigned to the d edges and $\mathbf{v}^\top \boldsymbol{\ell}_t$ is the total loss assigned to the path \mathbf{v} . Another (non-essential) simplifying assumption is that every action $\mathbf{v} \in \mathcal{S}$ has the same number of 1's: $\|\mathbf{v}\|_1 = m$ for all $\mathbf{v} \in \mathcal{S}$. The value of m plays an important role in the bounds below.

The proposed prediction algorithm is defined as follows. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be independent Gaussian random vectors taking values in \mathbb{R}^d such that the components of each \mathbf{X}_t are i.i.d. normal $X_{i,t} \sim \mathcal{N}(0, \eta^2)$ for some fixed $\eta > 0$ whose value will be specified later. Denote

$$\mathbf{Z}_t = \sum_{s=1}^t \mathbf{X}_s.$$

At time t , the forecaster chooses the action

$$\mathbf{V}_t = \arg \min_{\mathbf{v} \in \mathcal{S}} \left\{ \mathbf{v}^\top (\mathbf{L}_{t-1} + \mathbf{Z}_t) \right\},$$

where $\mathbf{L}_t = \sum_{s=1}^t \boldsymbol{\ell}_s$ for $t \geq 1$ and $\mathbf{L}_0 = (0, \dots, 0)^\top$.

The next theorem bounds the performance of the proposed forecaster. Again, we are not only interested in the regret but also the number of switches $\sum_{t=1}^n \mathbb{I}\{\mathbf{V}_{t+1} \neq \mathbf{V}_t\}$. The regret is of similar order—roughly $m\sqrt{dn}$ —as that of the standard FPL forecaster, up to a logarithmic factor. Moreover, the expected number of switches is $O(m \log d \sqrt{n})$. Remarkably, the dependence on d is only logarithmic and it is the weight m of the actions that plays an important role.

We note in passing that the Shrinking Dartboard algorithm of [Geulen et al. \(2010\)](#) can be used for simultaneously guaranteeing that the expected regret is $O(m^{3/2} \sqrt{n \log d})$ and the expected number of switches is $\sqrt{mn \log d}$. However, as this algorithm requires explicit computation of the exponential weighted forecaster, it can only be efficiently implemented for some special decision sets \mathcal{S} —see [Koolen et al. \(2010\)](#) and [Cesa-Bianchi and Lugosi \(2012\)](#) for some examples. On the other hand, our algorithm can be efficiently implemented whenever there exists an efficient implementation of the static optimization problem of finding $\arg \min_{\mathbf{v} \in \mathcal{S}} \mathbf{v}^\top \boldsymbol{\ell}$ for any $\boldsymbol{\ell} \in \mathbb{R}^d$. We also note that using FPL with identical perturbations does not guarantee any upper bound on the number of switches in the combinatorial setting, since the proof of [Kalai and Vempala \(2005\)](#) crucially depends on using exponentially distributed perturbations for *each action*. Clearly, this is not achievable by using perturbation vectors with exponentially distributed components.

Theorem 4 *Fix any $\mathbf{v} \in \mathcal{S}$. The expected regret and the expected number of action switches satisfy (under the oblivious adversary model)*

$$\mathbb{E} \widehat{L}_n - \mathbf{v}^\top \mathbf{L}_n \leq m\sqrt{n} \left(\frac{2d}{\eta} + \eta\sqrt{2 \log d} \right) + \frac{md(\log n + 1)}{\eta^2}$$

and

$$\begin{aligned} \mathbb{E} \sum_{t=1}^n \mathbb{I}\{\mathbf{V}_{t+1} \neq \mathbf{V}_t\} &\leq \sum_{t=1}^n \frac{m(1 + 2\eta\sqrt{2\log d} + \eta^2(2\log d + \sqrt{2\log d} + 1))}{4\eta^2 t} \\ &\quad + \sum_{t=1}^n \frac{m(1 + \eta\sqrt{2\log d})\sqrt{2\log d}}{\eta\sqrt{t}}. \end{aligned}$$

In particular, setting $\eta = \sqrt{\frac{2d}{\sqrt{2\log d}}}$ yields

$$\mathbb{E}\widehat{L}_n - \mathbf{v}^\top \mathbf{L}_n \leq 4m\sqrt{dn} \sqrt[4]{\log d} + m(\log n + 1)\sqrt{\log d}.$$

and

$$\mathbb{E} \sum_{t=1}^n \mathbb{I}\{\mathbf{V}_{t+1} \neq \mathbf{V}_t\} = O(m \log d \sqrt{n}).$$

The proof of the regret bound is quite standard, similar to the proof of Theorem 3 in [Audibert et al. \(2012\)](#), and is thus omitted. The more interesting part is the bound for the expected number of action switches $\mathbb{E} \sum_{t=1}^n \mathbb{I}\{\mathbf{V}_{t+1} \neq \mathbf{V}_t\} = \sum_{t=1}^n \mathbb{P}[\mathbf{V}_{t+1} \neq \mathbf{V}_t]$. It follows from the lemma below and the well-known facts that $\mathbb{E}[\|\mathbf{X}_1\|_\infty] \leq \sqrt{2\log d}$ and $\mathbb{E}[\|\mathbf{X}_1\|_\infty^2] \leq 2\log d + \sqrt{2\log d} + 1$ (see, e.g., [Boucheron et al. \(2013\)](#)). Thus, it suffices to prove the following:

Lemma 5 For each $t = 1, 2, \dots, n$,

$$\mathbb{P}[\mathbf{V}_{t+1} \neq \mathbf{V}_t | \mathbf{X}_{t+1}] \leq \frac{m \|\boldsymbol{\ell}_t + \mathbf{X}_{t+1}\|_\infty^2}{2\eta^2 t} + \frac{m \|\boldsymbol{\ell}_t + \mathbf{X}_{t+1}\|_\infty \sqrt{2\log d}}{\eta\sqrt{t}}.$$

Following the proof of Lemma 2, it is very easy to prove the following alternative regret bound:

Theorem 6 Fix any $\mathbf{v} \in \mathcal{S}$. The expected regret satisfies (under the oblivious adversary model)

$$\mathbb{E}\widehat{L}_n - \mathbf{v}^\top \mathbf{L}_n \leq \sum_{t=2}^n \mathbb{E} \left[\frac{m^2 \|\boldsymbol{\ell}_t + \mathbf{X}_{t+1}\|_\infty^3}{2\eta^2 t} + \frac{m^2 \|\boldsymbol{\ell}_t + \mathbf{X}_{t+1}\|_\infty^2 \sqrt{2\log d}}{\eta\sqrt{t}} \right].$$

In particular, setting $\eta = \frac{1}{2\sqrt{\log d}}$ yields

$$\mathbb{E}\widehat{L}_n - \mathbf{v}^\top \mathbf{L}_n = O(m^2 \log d \sqrt{n})$$

and

$$\mathbb{E} \sum_{t=1}^n \mathbb{I}\{\mathbf{V}_{t+1} \neq \mathbf{V}_t\} = O(m \log d \sqrt{n}).$$

The proof is left as an exercise. Note that this is the first known regret bound for FPL that depends on d only logarithmically. In particular, Theorem 6 improves the guarantees given in Theorem 4 for decision sets where $m = o(\sqrt{d})$. We present the proof of Lemma 5 below. **Proof** We use the notation $\mathbb{P}_t[\cdot] = \mathbb{P}[\cdot | \mathbf{X}_{t+1}]$ and $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot | \mathbf{X}_{t+1}]$. Also, let

$$\mathbf{h}_t = \boldsymbol{\ell}_t + \mathbf{X}_{t+1} \quad \text{and} \quad \mathbf{H}_t = \sum_{s=0}^{t-1} \mathbf{h}_s.$$

Furthermore, we will use the shorthand notation $c = \|\mathbf{h}_t\|_\infty$. Define the set A_t as the lead pack:

$$A_t = \left\{ \mathbf{w} \in \mathcal{S} : (\mathbf{w} - \mathbf{V}_t)^\top \mathbf{H}_t \leq \|\mathbf{w} - \mathbf{V}_t\|_1 c \right\}.$$

Observe that the choice of c guarantees that no action outside A_t can take the lead at time $t + 1$, since if $\mathbf{w} \notin A_t$, then

$$(\mathbf{w} - \mathbf{V}_t)^\top \mathbf{H}_t > \left| (\mathbf{w} - \mathbf{V}_t)^\top \mathbf{h}_t \right|$$

so $\mathbf{w}^\top \mathbf{H}_{t+1} > \mathbf{V}_t^\top \mathbf{H}_{t+1}$ and \mathbf{w} cannot be the new leader. It follows that we can upper bound the probability of switching as

$$\mathbb{P}_t[\mathbf{V}_{t+1} \neq \mathbf{V}_t] \leq \mathbb{P}_t[|A_t| > 1],$$

which leaves us with the problem of upper bounding $\mathbb{P}_t[|A_t| > 1]$. Similarly to the proof of Lemma 3, we start analyzing $\mathbb{P}_t[|A_t| = 1]$:

$$\begin{aligned} \mathbb{P}_t[|A_t| = 1] &= \sum_{\mathbf{v} \in \mathcal{S}} \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : (\mathbf{w} - \mathbf{v})^\top \mathbf{H}_t > \|\mathbf{w} - \mathbf{v}\|_1 c \right] \\ &= \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} f_{\mathbf{v}}(y) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y + \|\mathbf{w} - \mathbf{v}\|_1 c \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy, \end{aligned} \quad (2)$$

where $f_{\mathbf{v}}$ is the distribution of $\mathbf{v}^\top \mathbf{H}_t$. Next we crucially use the fact that the conditional distributions of correlated Gaussian random variables are also Gaussian. In particular, defining $k(\mathbf{w}, \mathbf{v}) = (m - \|\mathbf{w} - \mathbf{v}\|_1)$, the covariances are given as

$$\text{cov} \left(\mathbf{w}^\top \mathbf{H}_t, \mathbf{v}^\top \mathbf{H}_t \right) = \eta^2 (m - \|\mathbf{w} - \mathbf{v}\|_1) t = \eta^2 k(\mathbf{w}, \mathbf{v}) t.$$

Let us organize all actions $\mathbf{w} \in \mathcal{S} \setminus \mathbf{v}$ into a matrix $\mathbf{W} = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{N-1})$. The conditional distribution of $\mathbf{W}^\top \mathbf{H}_t$ is an $(N - 1)$ -variate Gaussian distribution with mean

$$\mu_{\mathbf{v}}(y) = \left(\mathbf{w}_1^\top \mathbf{L}_{t-1} + y \frac{k(\mathbf{w}_1, \mathbf{v})}{m}, \mathbf{w}_2^\top \mathbf{L}_{t-1} + y \frac{k(\mathbf{w}_2, \mathbf{v})}{m}, \dots, \mathbf{w}_{N-1}^\top \mathbf{L}_{t-1} + y \frac{k(\mathbf{w}_{N-1}, \mathbf{v})}{m} \right)^\top$$

and covariance matrix $\Sigma_{\mathbf{v}}$, given that $\mathbf{v}^\top \mathbf{H}_t = y$. Defining $\mathbf{k}_{\mathbf{v}} = (k(\mathbf{w}_1, \mathbf{v}), \dots, k(\mathbf{w}_{N-1}, \mathbf{v}))^\top$ and using the notation $\varphi(x) = \frac{1}{\sqrt{(2\pi)^{N-1} |\Sigma_{\mathbf{v}}|}} \exp(-\frac{x^2}{2})$, we get that

$$\begin{aligned} & \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y + \|\mathbf{w} - \mathbf{v}\|_1 c \mid \mathbf{v}^\top \mathbf{H}_t = y \right] \\ &= \int_{z_i=y+(m-k(\mathbf{w}_i, \mathbf{v}))c}^{\infty} \cdots \int^{\infty} \varphi \left(\sqrt{(\mathbf{z} - \mu_{\mathbf{v}}(y))^\top \Sigma_{\mathbf{v}}^{-1} (\mathbf{z} - \mu_{\mathbf{v}}(y))} \right) dz \\ &= \int_{z_i=y+mc}^{\infty} \cdots \int^{\infty} \varphi \left(\sqrt{(\mathbf{z} - \mu_{\mathbf{v}}(y) - c\mathbf{k}_{\mathbf{v}})^\top \Sigma_{\mathbf{v}}^{-1} (\mathbf{z} - \mu_{\mathbf{v}}(y) - c\mathbf{k}_{\mathbf{v}})} \right) dz \\ &= \int_{z_i=y+mc}^{\infty} \cdots \int^{\infty} \varphi \left(\sqrt{(\mathbf{z} - \mu_{\mathbf{v}}(y+mc))^\top \Sigma_{\mathbf{v}}^{-1} (\mathbf{z} - \mu_{\mathbf{v}}(y+mc))} \right) dz \\ &= \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y + mc \mid \mathbf{v}^\top \mathbf{H}_t = y + mc \right], \end{aligned}$$

where we used $\mu_{\mathbf{v}}(y+mc) = \mu_{\mathbf{v}}(y) + c\mathbf{k}_{\mathbf{v}}$. Using this, we rewrite (2) as

$$\begin{aligned} \mathbb{P}_t [|A_t| = 1] &= \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} f_{\mathbf{v}}(y) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy \\ &\quad - \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} (f_{\mathbf{v}}(y) - f_{\mathbf{v}}(y - mc)) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy \\ &= 1 - \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} (f_{\mathbf{v}}(y) - f_{\mathbf{v}}(y - mc)) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy. \end{aligned}$$

To treat the remaining term, we use that $\mathbf{v}^\top \mathbf{H}_t$ is Gaussian with mean $\mathbf{v}^\top \mathbf{L}_{t-1}$ and variance $\eta^2 mt$ and obtain

$$\begin{aligned} f_{\mathbf{v}}(y) - f_{\mathbf{v}}(y - mc) &= f_{\mathbf{v}}(y) \left(1 - \frac{f_{\mathbf{v}}(y - mc)}{f_{\mathbf{v}}(y)} \right) \\ &\leq f_{\mathbf{v}}(y) \left(\frac{mc^2}{2\eta^2 t} - \frac{c(y - \mathbf{v}^\top \mathbf{L}_{t-1})}{\eta^2 t} \right). \end{aligned}$$

Thus,

$$\begin{aligned} \mathbb{P}_t [|A_t| > 1] &\leq \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} (f_{\mathbf{v}}(y) - f_{\mathbf{v}}(y - mc)) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy \\ &\leq \sum_{\mathbf{v} \in \mathcal{S}} \int_{y \in \mathbb{R}} f_{\mathbf{v}}(y) \left(\frac{mc^2}{2\eta^2 t} - \frac{c(y - \mathbf{v}^\top \mathbf{L}_{t-1})}{\eta^2 t} \right) \mathbb{P}_t \left[\forall \mathbf{w} \neq \mathbf{v} : \mathbf{w}^\top \mathbf{H}_t > y \mid \mathbf{v}^\top \mathbf{H}_t = y \right] dy \\ &= \frac{mc^2}{2\eta^2 t} - \frac{c \mathbb{E} [\mathbf{V}_t^\top \mathbf{Z}_t]}{\eta^2 t} \leq \frac{mc^2}{2\eta^2 t} + \frac{mc \mathbb{E} [\|\mathbf{Z}_t\|_\infty]}{\eta^2 t}. \end{aligned}$$

Using the definition of c and $\mathbb{E} [\|\mathbf{Z}_t\|_\infty] \leq \eta \sqrt{2t \log d}$ gives the result. \blacksquare

References

- J. Y. Audibert, S. Bubeck, and G. Lugosi. Minimax policies for combinatorial prediction games. In *Conference on Learning Theory*, 2011. URL <http://arxiv.org/abs/1105.4871>.
- J. Y. Audibert, S. Bubeck, and G. Lugosi. Regret in online combinatorial optimization. *Manuscript*, 2012. URL <http://arxiv.org/abs/1204.4710>.
- S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.
- N. Cesa-Bianchi and G. Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78:1404–1422, 2012.
- E. Even-Dar, S. Kakade, and Y. Mansour. Online Markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009. ISSN 0364-765X. doi: <http://dx.doi.org/10.1287/moor.1090.0396>.
- W. Feller. *An Introduction to Probability Theory and its Applications, Vol. 1*. John Wiley, New York, 1968.
- Y. Freund and R. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55:119–139, 1997.
- C. Gentile and M. Warmuth. Linear hinge loss and average margin. In *Advances in Neural Information Processing Systems (NIPS)*, pages 225–231, 1998.
- S. Geulen, B. Voeking, and M. Winkler. Regret minimization for online buffering problems using the weighted majority algorithm. In *Proceedings of the Twenty-Third Conference on Computational Learning Theory*, pages 132–143, 2010.
- A. Grove, N. Littlestone, and D. Schuurmans. General convergence results for linear discriminant updates. *Machine Learning*, 43:173–210, 2001.
- A. György and G. Neu. Near-optimal rates for limited-delay universal lossy source coding. In *Proceedings of the IEEE International Symposium on Information Theory (ISIT)*, pages 2344–2348, 2011.
- J. Hannan. Approximation to Bayes risk in repeated play. *Contributions to the theory of games*, 3:97–139, 1957.
- E. Hazan, S. Kale, and M. Warmuth. Learning rotations with little regret. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 144–154, 2010.
- D. P. Helmbold and M. Warmuth. Learning permutations with exponential weights. *Journal of Machine Learning Research*, 10:1705–1736, 2009.

- M. Hutter and J. Poland. Prediction with expert advice by following the perturbed leader for general weights. In *Algorithmic Learning Theory*, pages 279–293. Springer, 2004.
- A. Kalai and S. Vempala. Efficient algorithms for the online decision problem. In B. Schölkopf and M. Warmuth, editors, *Proceedings of the 16th Annual Conference on Learning Theory and the 7th Kernel Workshop, COLT-Kernel 2003*, pages 26–40, New York, USA, Aug. 2003. Springer.
- A. Kalai and S. Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71:291–307, 2005.
- J. Kivinen and M. Warmuth. Relative loss bounds for multidimensional regression problems. *Machine Learning*, 45:301–329, 2001.
- W. Koolen, M. Warmuth, and J. Kivinen. Hedging structured concepts. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 93–105, 2010.
- N. Littlestone and M. Warmuth. The weighted majority algorithm. *Information and Computation*, 108:212–261, 1994.
- G. Neu, A. Györfgy, Cs. Szepesvári, and A. Antos. Online Markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems 23*, 2010.
- J. Poland. FPL analysis for adaptive bandits. In *In 3rd Symposium on Stochastic Algorithms, Foundations and Applications (SAGA '05)*, pages 58–69, 2005.
- A. Rakhlin, O. Shamir, and K. Sridharan. Relax and randomize : From value to algorithms. In *Advances in Neural Information Processing Systems 25*, pages 2150–2158, 2012.
- E. Takimoto and M. Warmuth. Paths kernels and multiplicative updates. *Journal of Machine Learning Research*, 4:773–818, 2003.
- V. Vovk. Aggregating strategies. In *Proceedings of the third annual workshop on Computational learning theory (COLT)*, pages 371–386, 1990.
- M. Warmuth and D. Kuzmin. Randomized online pca algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9:2287–2320, 2008.