Multiplicative Weights Update

Mukul Agarwal

September 9, 2026

These notes were prepared for a presentation which I gave with Abhigyan Dutta for the class **CS 592: Machine Learning Algorithms** (taught by Prof. Paul Valiant).

We begin with the general algorithmic setup, then state and analyze the Hedge algorithm. After discussing an application (AdaBoost), we prove a lower bound for online algorithms and show that MWU is optimal in this sense. The basic inequalities used are collected in the appendix.

For the most part, these notes follow the proofs and exposition of Arora, Hazan, and Kale [1].

1 Introduction

Consider the problem of allocating a portfolio among n stocks over T rounds. In each round t:

1. Select an allocation $\mathbf{p}^{(t)} \in \Delta_n$ based on the costs observed in previous rounds.
2. The cost vector $\mathbf{m}^{(t)} \in [-1, 1]^n$ is revealed.
3. Pay the round-t cost

$$\mathbf{p}_1^{(t)}\mathbf{m}_1^{(t)} + \ldots + \mathbf{p}_n^{(t)}\mathbf{m}_n^{(t)} = \left\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)} \right\rangle.$$

We make no assumptions about how the stocks behave. The cost vectors can be seen as form a matrix $\mathbf{M} \in [-1, 1]^{T \times n}$ whose rows are revealed one at a time. For example,

$$\mathbf{M} = \begin{pmatrix} 1.0 & -0.5 & \cdots & 0.7 \\ -0.3 & 0.0 & \cdots & -1.0 \\ \vdots & \vdots & \ddots & \vdots \\ 0.4 & -0.8 & \cdots & 0.2 \end{pmatrix}.$$

So, it is hard for us to do "well" in general. Instead, our goal is to compete with the best single stock in *hindsight*, obtaining a bound of the form

$$\sum_{t\in[T]} \left\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)} \right\rangle \leq \left(\min_{i\in[n]} \sum_{t\in[T]} \mathbf{m}_i^{(t)} \right) + \mathbf{regret},$$

with regret as small as possible.

1.1. Notation

For convenience, we will write

$$\mathcal{L}_i^{(t)} = \sum_{\tau\in[t]} \mathbf{m}_i^{(\tau)} \quad \text{and} \quad \mathcal{L}_i \coloneqq \mathcal{L}_i^{(T)} = \sum_{\tau\in[T]} \mathbf{m}_i^{(\tau)}.$$

The things we refer to as "stocks" are generally called "experts" in the literature (though this is a bit confusing since we don't have any guarantees about how good the experts are). In the case of **AdaBoost**, we may also refer to them as examples, for reasons that will become clear.

2 Hedge

Let us discuss the Hedge algorithm, originally due to Yoav Freund and Robert Schapire, which is an example of the MWU framework.

Fix a learning rate $\eta \in (0,1)$, which parametrizes the algorithm. and . Then, the Hedge algorithm is as follows:

1. Initialize $\mathbf{w}^{(1)} \leftarrow \mathbf{1} \in \mathbb{R}^n_{>0}$ to be the all-ones (uniform) vector.
2. For rounds $t = 1, 2, \ldots, T$:
    1. Allocate to each "stock" in proportion to its weight, i.e., according to $\mathbf{p}^{(t)}$:

$$\mathbf{p}^{(t)} = \frac{\mathbf{w}^{(t)}}{\left\|\mathbf{w}^{(t)}\right\|_1} \in \Delta_n. \qquad \left(\text{so } \mathbf{p}_i^{(t)} = \frac{\mathbf{w}_i^{(t)}}{\sum_{i\in[n]} \mathbf{w}_i^{(t)}} \right)$$

    2. Observe returns $\mathbf{m}^{(t)}$.
    3. Update the weights as

$$\mathbf{w}_i^{(t+1)} \leftarrow \mathbf{w}_i^{(t)} \cdot \exp\left(-\eta \mathbf{m}_i^{(t)}\right) = \exp\left(-\eta \mathcal{L}_i^{(t)}\right).$$

We now analyze how competitive this algorithm is.

2.1. Single-Expert Analysis

Theorem 1 **Upper bound guarantee**

For every $\eta \in (0,1)$, $T \geq 1$, and sequence $\mathbf{m}^{(t)} \in [-1,1]^n$, Hedge satisfies

$$\sum_{t\in[T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \leq \min_i \mathcal{L}_i + \eta \sum_{t\in[T]} \left\langle \left(\mathbf{m}^{(t)}\right)^2, \mathbf{p}^{(t)} \right\rangle + \frac{\log n}{\eta}.$$

Define the potential $\Phi^{(t)}$ to be the ℓ_1-norm of $\mathbf{w}^{(t)}$, i.e.,

$$\Phi^{(t)} = \sum_{i\in[n]} \mathbf{w}_i^{(t)}$$

Lower Bounding $\Phi^{(t+1)}$. First, we lower bound the potential. Put

$$i^* = \operatorname*{argmin}_{i\in[n]} \mathcal{L}_i$$

to be the best "expert"/"stock". Then, we clearly have that

$$\Phi^{(T+1)} \geq \mathbf{w}_{i^*}^{(T+1)} = \prod_{\tau\in[T]} \exp\left(-\eta \mathbf{m}_i^{(\tau)}\right) = \exp(-\eta \mathcal{L}_i).$$

Upper Bounding $\Phi^{(t+1)}$. We know that

$$\begin{aligned}
\Phi^{(t)} &= \sum_{i\in[n]} \mathbf{w}_i^{(t+1)} \\
&= \sum_{i\in[n]} \mathbf{w}_i^{(t)} \exp\left(-\eta \mathbf{m}_i^{(t)}\right) \\
&= \sum_{i\in[n]} \Phi^{(t)} \mathbf{p}_i^{(t)} \exp\left(-\eta \mathbf{m}_i^{(t)}\right) \\
&= \Phi^{(t)} \sum_{i\in[n]} \mathbf{p}_i^{(t)} \exp\left(-\eta \mathbf{m}_i^{(t)}\right) \\
&\overset{(a)}{\leq} \Phi(t) \sum_{i\in[n]} \mathbf{p}_i^{(t)} \left(1 - \eta \mathbf{m}_i^{(t)} + \eta \left(\mathbf{m}_i^{(t)}\right)^2\right) \\
&= \Phi^{(t)} \left[\sum_{i\in[n]} \mathbf{p}_i^{(t)} - \eta \sum_{i\in[n]} \mathbf{p}_i^{(t)} \mathbf{m}_i^{(t)} + \eta^2 \sum_{i\in[n]} \mathbf{p}_i^{(t)} \left(\mathbf{m}_i^{(t)}\right)^2\right] \\
&\overset{(b)}{=} \Phi^{(t)} \left[1 - \eta \left\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)} \right\rangle + \eta^2 \left\langle \mathbf{p}^{(t)}, \left(\mathbf{m}^{(t)}\right)^2 \right\rangle\right]
\end{aligned}$$

$$\overset{(c)}{\leq} \Phi^{(t)} \exp\Big(-\eta\big\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)}\big\rangle + \eta^2 \Big\langle \mathbf{p}^{(t)}, \big(\mathbf{m}^{(t)}\big)^2 \Big\rangle\Big).$$

Here (a) is by $e^x \leq 1 + x + x^2$, (b) is because $\mathbf{p}^{(t)} \in \Delta_n$ so $\sum_{i\in[n]} \mathbf{p}_i^{(t)} = 1$, and (c) is by the inequality $1 + x \leq e^x$.

Applying this inequality recursively and going from $\Phi^{(T+1)}$ to $\Phi^{(1)} = n$ we obtain that

$$\begin{aligned}\Phi^{(T+1)} &\leq n \prod_{t\in[T]} \exp\Big(-\eta\big\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)}\big\rangle + \eta^2 \Big\langle \mathbf{p}^{(t)}, \big(\mathbf{m}^{(t)}\big)^2 \Big\rangle\Big) \\ &= n \exp\left(-\eta \sum_{t\in[T]} \big\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)}\big\rangle + \eta^2 \sum_{t\in[T]} \Big\langle \mathbf{p}^{(t)}, \big(\mathbf{m}^{(t)}\big)^2 \Big\rangle\right).\end{aligned}$$

Combining the bounds. Using transitivity of $(\leq)$ we get an inequality independent of $\Phi^{(t)}$. Taking logarithms on both sides, we obtain

$$-\eta \mathcal{L}_{i^*} \leq \log(n) - \eta \sum_{t\in[T]} \big\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)}\big\rangle + \eta^2 \sum_{t\in[T]} \Big\langle \mathbf{p}^{(t)}, \big(\mathbf{m}^{(t)}\big)^2 \Big\rangle.$$

Dividing by η and rearranging terms, we get

$$\sum_{t\in[T]} \big\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)}\big\rangle \leq \min_i \mathcal{L}_i + \eta \sum_{t\in[T]} \Big\langle \big(\mathbf{m}^{(t)}\big)^2, \mathbf{p}^{(t)} \Big\rangle + \frac{\log n}{\eta},$$

as desired.

2.2. Set-of-experts analysis

The analysis can be refined by instead considering a *set* of experts to compete against; this is the form which will be used when analyzing **AdaBoost**. Note that the previous analysis can be recovered as the special case where $E = \{i^*\}$.

Theorem 2 **Upper bound guarantee**

For every $\eta \in (0, 1)$, $T \geq 1$, subset $E \subseteq [n]$ and sequence $\mathbf{m}^{(t)} \in [-1, 1]^n$, Hedge satisfies

$$\sum_{t\in[T]} \big\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)}\big\rangle \leq \frac{1}{|E|} \sum_{i\in E} \mathcal{L}_i + \eta \sum_{t\in[T]} \Big\langle \big(\mathbf{m}^{(t)}\big)^2, \mathbf{p}^{(t)} \Big\rangle + \frac{1}{\eta} \log\left(\frac{n}{|E|}\right).$$

Notice that for general E we are giving something up in the first term but are gaining in the $\log(n)$ term. In order to prove this, we can use the same upper bound on the potential, but a slightly more refined lower bound.

Improved Lower Bound on $\Phi^{(T+1)}$. We have that

$$\begin{aligned}
\Phi^{(T+1)} &\geq \sum_{i \in E} \mathbf{w}_i^{(T+1)} \\
&= \sum_{i \in E} \exp(-\eta \mathcal{L}_i) \\
&= |E| \cdot \underbrace{\frac{1}{|E|} \sum_{i \in E} \exp(-\eta \mathcal{L}_i)}_{\text{Average of } \exp(\cdot) \text{ on } \{\mathcal{L}_i\}_{i \in E}} \\
&\geq |E| \cdot \exp\left(\frac{1}{|E|} \sum_{i \in E} (-\eta \mathcal{L}_i) \right),
\end{aligned}$$

where the last inequality is due to Jensen's inequality on $\exp(\cdot)$.

Combining bounds. Again using transitivity and taking a logarithm, we get

$$-\log|E| - \frac{\eta}{|E|} \sum_{i \in E} \mathcal{L}_i \leq \log(n) - \eta \sum_{t \in [T]} \left\langle \mathbf{p}^{(t)}, \mathbf{m}^{(t)} \right\rangle + \eta^2 \sum_{t \in [T]} \left\langle \mathbf{p}^{(t)}, \left(\mathbf{m}^{(t)}\right)^2 \right\rangle.$$

Rearranging terms and dividing by η, we get the desired result

$$\sum_{t \in [T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \leq \frac{1}{|E|} \sum_{i \in E} \mathcal{L}_i + \eta \sum_{t \in [T]} \left\langle \left(\mathbf{m}^{(t)}\right)^2, \mathbf{p}^{(t)} \right\rangle + \frac{1}{\eta} \log\left(\frac{n}{|E|} \right).$$

3 Applications

We now discuss some applications. The first amongst these is Boosting the accuracy of a Learner. To discuss this, we first define the idea of PAC learning, first proposed by Valiant in 1984.

3.1. PAC Learning

Let $\mathcal{X}$ be a set of examples. A concept $c : \mathcal{X} \to \{0, 1\}$ assigns a label to each example, and a concept class $\mathcal{C}$ is a set of such functions. A hypothesis $h : \mathcal{X} \to \{0, 1\}$ is a proposed classifier, drawn from a hypothesis class $\mathcal{H}$. Under a distribution $\mathcal{D}$ on $\mathcal{X}$, its error with respect to c is

$$\mathrm{err}_{\mathcal{D}}(h, c) := \Pr_{x \sim \mathcal{D}}[h(x) \neq c(x)].$$

Definition 3.1

A concept class $\mathcal{C}$ is (strongly) PAC-learnable if there is an algorithm $\mathcal{A}$ that, for every $c \in \mathcal{C}$, distribution $\mathcal{D}$ on $\mathcal{X}$, and $\varepsilon, \delta > 0$, uses labeled examples from $\mathcal{D}$ to find a hypothesis $h \in \mathcal{H}$ satisfying

$$\mathrm{err}_{\mathcal{D}}(h, c) \leq \varepsilon$$

with probability at least $1 - \delta$.[1]

A weak learner instead guarantees error at most $1/2 - \gamma$ for some $\gamma > 0$ (this value γ can be called the **advantage**). Boosting uses this learner repeatedly to obtain a hypothesis with error below an arbitrary $\varepsilon > 0$.

3.2. Boosting on a Fixed Set of Examples

Fix n examples, indexed by $[n]$, with labels given by $c : [n] \to \{0, 1\}$. Our goal is to produce a hypothesis h^* that misclassifies fewer than an ε fraction of these examples:

$$\mathrm{err}_{\mathcal{U}(n)}(h^*, c) = \frac{\left|\left\{i \in [n] : h^{*(i)} \neq c(i)\right\}\right|}{n} < \varepsilon,$$

where $\mathcal{U}(n)$ is the uniform distribution on $[n]$.

Suppose the weak learner $\mathcal{A}$ can achieve error at most $1/2 - \gamma$ under any distribution on these examples. We will change this distribution between rounds, run $\mathcal{A}$ repeatedly, and combine the resulting hypotheses by majority vote. The examples themselves stay fixed; Hedge maintains a weight for each example.

3.3. AdaBoost

Fix parameters η and T, to be chosen in the analysis. The algorithm is an insantiation of **Hedge**, as follows:

1. Initialize $\mathbf{w}^{(1)} \leftarrow \mathbf{1} \in \mathbb{R}^n_{>0}$ to be the all-ones (uniform) vector.
2. For rounds $t = 1, 2, \ldots, T$:
    1. Run the weak learner $\mathcal{A}$ with distribution $\mathbf{p}^{(t)} = \mathbf{w}^{(t)} / \left\|\mathbf{w}^{(t)}\right\|_1$ to obtain a hypothesis $h^{(t)}$
    2. Let $\mathbf{m}^{(t)}$ indicate which examples are classified correctly:

[1] The usual definition also requires polynomial bounds on the number of examples and the runtime.

$$\mathbf{m}_i^{(t)} \leftarrow \begin{cases} 0 \text{ if } h^{(t)}(i) \neq c(i) \\ 1 \text{ if } h^{(t)}(i) = c(i) \end{cases}.$$

3. Update the weights as

$$\mathbf{w}_i^{(t+1)} \leftarrow \mathbf{w}_i^{(t)} \cdot \exp\left(-\eta \mathbf{m}_i^{(t)}\right) \quad \text{for every } i \in [n].$$

3. Return the final hypothesis to be the majority vote:

$$h^{*(i)} \leftarrow \operatorname*{maj}_{t \in [T]} h^{(t)}(i).$$

Correctly classified examples lose weight, so the next round places relatively more weight on the examples the learner misclassified. We now choose η and T so that the final hypothesis has error below ε with high probability.

3.4. Analysis

Theorem 2 **Finite-domain boosting**

Assume that for any distribution $\mathcal{D}$ on $[n]$ our weak learners attain error $\leq 1/2 - \gamma$ with probability at least $1 - \delta'$. If we wish to attain

$$\operatorname{err}_{\mathcal{U}(n)}(h^*, c) < \varepsilon$$

with probability at least $1 - \delta$, then it suffices to choose parameters

$$\eta = \frac{\gamma}{2}, \qquad T > \frac{4}{\gamma^2} \log\left(\frac{1}{\varepsilon}\right), \quad \text{and} \quad \delta' = \frac{\delta}{T} = \frac{\gamma^2 \delta}{4 \log(1/\varepsilon)}.$$

The general idea of the proof is to choose a set E to be the set on which our final hypothesis errs, and then using the weak learning assumption along with the subset bound on Hedge to prove that E is sufficiently small.

Using the Weak-Learning Assumption. We know that an individual weak learner is erroneous on at most $1/2 - \gamma$ of the samples, so it is correct on at least $\frac{1}{2} + \gamma$ of the samples. Since the probability with respect to the distribution at the t-th step is given by $\mathbf{p}^{(t)}$ and the vector $\mathbf{m}^{(t)}$ is an indicator of correct classification,

$$\left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle = \Pr_{x \sim \mathcal{D}[\mathbf{p}^{(t)}]} \left[h^{(t)}(x) = c(x) \right] \geq \frac{1}{2} + \gamma.$$

Summing over the T rounds, we immediately get

$$\sum_{t\in[T]} \langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \rangle \geq \left(\frac{1}{2} + \gamma\right) T.$$

But this is the left-hand side of the Hedge bound! So we can lower bound the right-hand side of the hedge bound by $(1/2 + \gamma)T$.

Choosing E. Let E be the set of elements misclassified by the final classifier:

$$E = \{i : h^*(i) \neq c(i)\}.$$

If E is empty, we see that h^* is zero-error, so we are done! Otherwise, we know that for any $i \in E$, i has to be correctly classified by *at most half* of the classifiers for it to be misclassified by h^*, i.e.,

$$\mathcal{L}_i = \sum_{t\in[T]} m_i \leq \frac{T}{2}.$$

We apply the subset-bound version of the Hedge upper bound to see

$$\begin{aligned}
\left(\frac{1}{2} + \gamma\right) T &\leq \sum_{t\in[T]} \langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \rangle \\
&\leq \frac{1}{|E|} \sum_{\substack{i\in E \\ t\in[T]}} \mathcal{L}_i + \eta \sum_{t\in[T]} \left\langle \left(\mathbf{m}^{(t)}\right)^2, \mathbf{p}^{(t)} \right\rangle + \frac{1}{\eta} \log\left(\frac{n}{|E|}\right) \\
&\leq \frac{T}{2} + \eta T + \frac{1}{\eta} \log\left(\frac{n}{|E|}\right),
\end{aligned}$$

and so we have

$$\log\left(\frac{n}{|E|}\right) \geq \eta(\gamma - \eta)T \Longrightarrow \log\left(\frac{|E|}{n}\right) \leq -\eta(\gamma - \eta)T.$$

If we choose $\eta = \frac{\gamma}{2}$ and exponentiate, then we get

$$\mathrm{err}_{\mathcal{U}(n)}(h^*, c) = \frac{|E|}{n} \leq \exp\left(-\frac{\gamma^2}{4} T\right).$$

So to drive the error below ε it suffices to choose

$$T > \frac{4}{\gamma^2} \log\left(\frac{1}{\varepsilon}\right).$$

Let the desired overall upper bound on failure probability be δ. Let F_t be the event that the t-th hypothesis fails to have error $\leq 1/2 - \gamma$. If we require that each weak

PAC learning iteration has probability of error at most δ/T, then by the union bound

$$\Pr[\text{boosting fails}] = \Pr\left[\bigcup_{t\in[T]} F_t\right] \le \sum_{t\in[T]} \Pr[F_t] \le \delta.$$

4 Lower Bound on Regret

Lemma 3 **Balancing Terms in Upper Bound**

For $n \ge 2$ and $T > \log n$, choosing $\eta = \sqrt{(\log n)/T}$ in Hedge guarantees

$$\sum_{t\in[T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \le \min_{i\in[n]} \mathcal{L}_i + 2\sqrt{T \log n}.$$

Proof: Since $\left\langle \left(\mathbf{m}^{(t)}\right)^2, \mathbf{p}^{(t)} \right\rangle \le 1$, the Hedge bound gives

$$\sum_{t\in[T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \le \min_{i\in[n]} \mathcal{L}_i + \eta T + \frac{\log n}{\eta}.$$

Balancing the two regret terms, we find that equality is attained when $\eta = \sqrt{\log n/T}$ (this is meaningful since $T > \log n$). Substituting this choice gives

$$\sum_{t\in[T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \le \min_{i\in[n]} \mathcal{L}_i + 2\sqrt{T \log n}. \qquad \square$$

The following lower bound shows that this regret is optimal up to an absolute constant.

Theorem 4 **Regret lower bound**

For every deterministic online allocation algorithm with $n \ge 8$ experts and even $T \ge C \log n$, for a sufficiently large absolute constant C, there exists a sequence $\mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(T)} \in \{0, 1/2, 1\}^n$ such that

$$\min_{i\in[n]} \sum_{t\in[T]} \mathbf{m}_i^{(t)} = \Omega(T),$$

and its allocations $\mathbf{p}^{(1)}, \ldots, \mathbf{p}^{(T)} \in \Delta_n$ satisfy

$$\sum_{t\in[T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \ge \min_{i\in[n]} \sum_{t\in[T]} \mathbf{m}_i^{(t)} + \Omega\left(\sqrt{T \log n}\right).$$

After fixing some algorithm, we construct the costs randomly and get a bound that holds with nonzero probability. After this, we apply the probabilistic method to get an existential guarantee that *some* sequence of vectors must attain such a bound.

The construction is as follows. Set $\mathbf{m}_1^{(t)} = 1/2$ for every $t \in [T]$; otherwise,

$$\mathbf{m}_i^{(t)} \sim \text{Bernoulli}\left(\frac{1}{2}\right) \quad \text{for every } i > 1 \text{ and } t \in [T],$$

and all random samples are independent. Since the algorithm is deterministic and the only information revealed is given in the previous costs, the deterministic algorithm's allocation $\mathbf{p}^{(t)}$ is fixed conditioned on the previous costs. Further, $\mathbb{E}\left[\mathbf{m}_i^{(t)}\right] = 1/2$ for each i, so

$$\mathbb{E}\left[\left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \mid \mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(t-1)}\right] = \sum_{i \in [n]} \mathbf{p}_i^{(t)} \cdot \frac{1}{2} = \frac{1}{2},$$

since $\mathbf{p}^{(t)} \in \Delta_n$. By the law of total expectation, each round also has unconditional expected cost $1/2$:

$$\begin{aligned}\mathbb{E}\left[\left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle\right] &= \mathbb{E}\left[\mathbb{E}\left[\left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle \mid \mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(t-1)}\right]\right] \\ &= \mathbb{E}\left[\frac{1}{2}\right] = \frac{1}{2},\end{aligned}$$

and linearity of expectation now gives

$$\mathbb{E}\left[\sum_{t \in [T]} \left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle\right] = \sum_{t \in [T]} \mathbb{E}\left[\left\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \right\rangle\right] = \frac{T}{2}.$$

Lemma 5 **Expected cost of the best expert**

For the random costs constructed above,

$$\mathbb{E}\left[\min_{i \in [n]} \mathcal{L}_i\right] \leq \frac{T}{2} - \Omega\left(\sqrt{T \log n}\right).$$

Proof: Set $u = 1/4\sqrt{T \log n}$. Since $T \geq C \log n$ for sufficiently large C, we have $u \leq T/8$. For each $i > 1$, the binomial lower-tail bound gives

$$\Pr\left[\mathcal{L}_i \leq \frac{T}{2} - u\right] \geq \frac{1}{15} \exp\left(-\frac{16u^2}{T}\right) = \frac{1}{15n}.$$

By independence and $1 + x \leq \exp(x)$,

$$\Pr\Big[\min_{i\in[n]} \mathcal{L}_i > \frac{T}{2} - u\Big] \le \Big(1 - \frac{1}{15n}\Big)^{n-1} \le \exp\Big(-\frac{n-1}{15n}\Big).$$

Thus, with probability $\Omega(1)$, some expert costs at most $T/2 - u$. The best expert always costs at most $T/2$, since expert 1 does. Therefore,

$$\mathbb{E}\Big[\min_{i\in[n]} \mathcal{L}_i\Big] \le \frac{T}{2} - \Omega(u) = \frac{T}{2} - \Omega\Big(\sqrt{T\log n}\Big). \qquad \square$$

Lemma 6 **Positive probability of large regret**

For the random costs constructed above, write

$$R := \sum_{t\in[T]} \big\langle \mathbf{m}^{(t)}, \mathbf{p}^{(t)} \big\rangle - \min_{i\in[n]} \mathcal{L}_i.$$

There is an absolute constant $c > 0$ such that

$$\Pr\Big[R \ge c\sqrt{T\log n}\Big] \ge \Omega\left(\sqrt{\frac{\log n}{T}}\right).$$

Proof: Put $\mu = \mathbb{E}[R]$. The expected-cost lemma and the algorithm's expected cost $T/2$ give $\mu \ge \Omega\big(\sqrt{T\log n}\big)$. Since $R \le T$, apply Markov's inequality to $T - R$:

$$\begin{aligned}
\Pr\Big[R \ge \frac{\mu}{2}\Big] &= 1 - \Pr\Big[R < \frac{\mu}{2}\Big] \\
&= 1 - \Pr\Big[T - R > T - \frac{\mu}{2}\Big] \\
&\ge 1 - \frac{T-\mu}{T - \mu/2} = \frac{\mu}{2T} = \Omega\left(\sqrt{\frac{\log n}{T}}\right).
\end{aligned}$$

Since $\mu = \Omega\big(\sqrt{T\log n}\big)$, this proves the claim. $\square$

We can now prove the main result. Let

$$A := \text{the event that } R \ge c\sqrt{T\log n}, \text{ and}$$
$$B := \text{the event that } \min_{i\in[n]} \mathcal{L}_i < \frac{T}{4}.$$

The large-regret lemma gives $\Pr[A] \ge d\sqrt{(\log n)/T}$ for an absolute constant $d > 0$.

For each $i > 1$, Hoeffding's inequality gives $\Pr[\mathcal{L}_i < T/4] \leq \exp(-T/8)$. Since $\mathcal{L}_1 = T/2$, the union bound gives

$$\Pr[B] \leq \sum_{i=2}^{n} \Pr\left[\mathcal{L}_i < \frac{T}{4}\right] \leq n \exp\left(-\frac{T}{8}\right).$$

Choose C sufficiently large. Since $T \geq C \log n$,

$$n \exp\left(-\frac{T}{8}\right) \leq \exp\left(-\frac{T}{16}\right) < d\sqrt{\frac{\log n}{T}}.$$

Both inequalities hold for sufficiently large C, since $\log n \leq T/C$ and $\sqrt{T}\exp(-T/16)$ tends to zero.

Applying the union bound once more,

$$\Pr[A \cap B^c] \geq \Pr[A] - \Pr[B] \geq d\sqrt{\frac{\log n}{T}} - n \exp\left(-\frac{T}{8}\right) > 0.$$

Hence, by the probabilistic method, there exist some cost vectors $\mathbf{m}^{(1)}, \ldots, \mathbf{m}^{(n)} \in \left\{0, \frac{1}{2}, 1\right\}^n$ which simultaneously satisfy

$$\frac{T}{4} \leq \min_{i \in [n]} \mathcal{L}_i \leq \frac{T}{2} \quad \text{and} \quad R \geq c\sqrt{T \log n},$$

completing the proof of the theorem.

5 Appendix: Basic Inequalities

Lemma 7 **Exponential bounds**

For every $x \in \mathbb{R}$, $1 + x \leq \exp(x)$. Moreover, for every $x \in [-1, 1]$,

$$\exp(x) \leq 1 + x + x^2.$$

Proof of Lower Bound: Let $f(x) = \exp(x)$. At $x = 0$, equality holds; for $x \neq 0$, Taylor's theorem at 0 with Lagrange remainder gives a point c between 0 and x such that

$$f(x) = 1 + x + \frac{\exp(c)}{2}x^2 \geq 1 + x. \qquad \square$$

Proof of Upper Bound: Let $f(x) = \exp(x)$ and $x \in [-1, 1]$. The case $x = 0$ is immediate. Otherwise, Taylor's theorem at 0 gives a point c between 0 and x such that

$$\begin{aligned} f(x) &= 1 + x + \frac{x^2}{2} + \frac{\exp(c)}{6}x^3 \\ &\leq 1 + x + \frac{x^2}{2} + \frac{|x^3|}{6} \cdot \left| \sup_{c \in [-1,1]} \exp(c) \right| \\ &\leq 1 + x + \frac{x^2}{2} + \left(\frac{x^2}{6}\right) \cdot 3 \leq 1 + x + x^2, \end{aligned}$$

since $\exp(c) \leq e < 3$ and $|x|^3 \leq x^2$ on $[-1, 1]$. $\square$

Lemma 8 **Jensen's inequality**

Let $I \subseteq \mathbb{R}^d$ be convex and let $f : I \to \mathbb{R}$ be convex. Take points $\mathbf{x}_1, \ldots, \mathbf{x}_n \in I$ and let $\lambda = (\lambda_1 \ \cdots \ \lambda_n) \in \Delta_n$. Then

$$f\left(\sum_{i=1}^{n} \lambda_i \mathbf{x}_i\right) \leq \sum_{i=1}^{n} \lambda_i f(\mathbf{x}_i).$$

Proof: Induct on n. The case $n = 1$ is equality. For $n \geq 2$, set $\alpha = 1 - \lambda_n$. If $\alpha = 0$, the claim is again equality. Otherwise, put $\mathbf{y} = \sum_{i \in [n-1]} (\lambda_i / \alpha)\mathbf{x}_i \in I$. By convexity and the induction hypothesis,

$$f\left(\sum_{i=1}^{n}\lambda_i \mathbf{x}_i\right) = f(\alpha \mathbf{y} + \lambda_n \mathbf{x}_n) \leq \alpha f(\mathbf{y}) + \lambda_n f(\mathbf{x}_n)$$

$$\leq \alpha \sum_{i=1}^{n-1} \frac{\lambda_i}{\alpha} f(\mathbf{x}_i) + \lambda_n f(\mathbf{x}_n).$$

$$= \sum_{i=1}^{n} \lambda_i \, f(\mathbf{x}_i). \qquad \square$$

We do not prove the following two inequalities, but direct the reader to **The Probabilistic Method** by Noga Alon and Joel Spencer, where more detailed proofs are written.

Lemma 9 **Hoeffding's inequality**

Let $X_1, ..., X_T$ be independent random variables in $[0, 1]$, and put $S = \sum_{t\in[T]} X_t$. For every $u \geq 0$,

$$\Pr[S \leq \mathbb{E}[S] - u] \leq \exp\left(-\frac{2u^2}{T}\right).$$

Proof idea: Use independence to write the MGF of $S - \mathbb{E}[S]$ as a product of the individual MGFs, and then use boundedness to bound each factor above by $\exp(O(s^2))$. Optimizing over s gives the desired tail bound. $\square$

Lemma 10 **Binomial lower tail**

If $X \sim \text{Bin}(T, 1/2)$ where T is a positive even integer, and $0 \leq u \leq T/8$, then

$$\Pr\left[X \leq \frac{T}{2} - u\right] \geq \frac{1}{15}\exp\left(-\frac{16u^2}{T}\right).$$

Proof Idea: Choose consecutive integers $k \leq T/2 - u$, and then bound each $\Pr[X = k]$ below by comparing it to $\Pr[X = T/2]$. Summing over these integers gives the desired lower bound. $\square$

References

[1] S. Arora, E. Hazan, and S. Kale, “The Multiplicative Weights Update Method: a Meta-Algorithm and Applications,” *Theory of Computing*, vol. 8, no. 6, pp. 121–164, May 2012, doi: 10.4086/toc.2012.v008a006.