Lecture 25: Online Learning

Michael Dinitz

December 2, 2025
601.433/633 Introduction to Algorithms
Slides by Michael Dinitz and Jessica Sorrell

Introduction

Machine Learning from the point of view of theoretical computer science

- Proofs about performance
- Minimize assumptions
- *Not* going to talk about useful in practice, etc.

Today:

- Online Learning

Online Learning

Online Learning

Learning over time, not just one-shot

- See data one piece at a time
- Try to use historical data to make decisions as we go
- We don't assume data comes from a distribution. Could be adversarially chosen!

Example: Learning From Expert Advice

Intuition: stock market

- N experts
- Every day:
  - Every expert predicts up/down
  - Algorithm makes prediction
  - Find out what happened

What can/should we do? Can we always make an accurate prediction?

Example: Learning From Expert Advice

Intuition: stock market

- N experts
- Every day:
  - Every expert predicts up/down
  - Algorithm makes prediction
  - Find out what happened

What can/should we do? Can we always make an accurate prediction?

- No! Experts could all be essentially random, uncorrelated with market

Example: Learning From Expert Advice

Intuition: stock market

- N experts
- Every day:
  - Every expert predicts up/down
  - Algorithm makes prediction
  - Find out what happened

What can/should we do? Can we always make an accurate prediction?

- No! Experts could all be essentially random, uncorrelated with market

Easier (but still interesting) goal: can we do as well as the best expert?

- Don't try to learn the market: learn which expert knows the market best

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Each day:

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Each day:

- Majority vote of remaining experts
- Remove incorrect experts

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Each day:

- Majority vote of remaining experts
- Remove incorrect experts

Best expert makes $\mathbf{0}$ mistakes

We make:

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Each day:

- Majority vote of remaining experts
- Remove incorrect experts

Best expert makes $\mathbf{0}$ mistakes

We make: $\boldsymbol{O(\log N)}$ mistakes

Warmup

Assume best expert makes $\mathbf{0}$ mistakes: always correctly predicts the market.
How should we predict market to minimize #mistakes?

Each day:

- Majority vote of remaining experts
- Remove incorrect experts

Best expert makes $\mathbf{0}$ mistakes

We make: $\boldsymbol{O(\log N)}$ mistakes

- Each mistake decreases # experts by $\mathbf{1/2}$

General case: no perfect expert

General case: no perfect expert

Weighted Majority

- Initialize all experts to weight $\mathbf{1}$
- Predict based on *weighted* majority vote
- Penalize mistakes by cutting weights in half

$\boldsymbol{M} = \#$ mistakes we've made
$\boldsymbol{m} = \#$ mistakes best expert has made
$\boldsymbol{W} = $ total weight

General case: no perfect expert

Weighted Majority

- Initialize all experts to weight $\mathbf{1}$
- Predict based on *weighted* majority vote
- Penalize mistakes by cutting weights in half

$\boldsymbol{M}$ = # mistakes we've made
$\boldsymbol{m}$ = # mistakes best expert has made
$\boldsymbol{W}$ = total weight

$\boldsymbol{W \geq (1/2)^m}$

- Best expert has weight at least $\boldsymbol{(1/2)^m}$

General case: no perfect expert

Weighted Majority

- Initialize all experts to weight 1
- Predict based on *weighted* majority vote
- Penalize mistakes by cutting weights in half

M = # mistakes we've made
m = # mistakes best expert has made
W = total weight

$W \geq (1/2)^m$

- Best expert has weight at least $(1/2)^m$

$W \leq N(3/4)^M$

- Every time we make a mistake, at least $1/2$ the total weight gets decreased by $1/2$, so left with at most $3/4$ of the original total weight

General case: no perfect expert

Weighted Majority

- Initialize all experts to weight $\mathbf{1}$
- Predict based on *weighted* majority vote
- Penalize mistakes by cutting weights in half

M = # mistakes we've made
m = # mistakes best expert has made
W = total weight

$W \geq (1/2)^m$

- Best expert has weight at least $(1/2)^m$

$W \leq N(3/4)^M$

- Every time we make a mistake, at least $1/2$ the total weight gets decreased by $1/2$, so left with at most $3/4$ of the original total weight

$$\implies (1/2)^m \leq W \leq N(3/4)^M \implies (4/3)^M \leq N2^m$$

$$\implies M \leq \log_{4/3}(N2^m) = \frac{m + \log N}{\log(4/3)} \approx 2.4(m + \log N)$$

A Better Algorithm (and More General Framework!)

What if we have more than two choices? What if some mistakes are “worse” than others?

A Better Algorithm (and More General Framework!)

What if we have more than two choices? What if some mistakes are “worse” than others?

General setup:

- T time steps (days)
- N actions the algorithm can take (experts)
- At each time step $t \in [T]$, algorithm A chooses an action $i \in [N]$
- Each action $i \in [N]$ then receives a loss $\ell_i^t \in [0, 1]$, and the algorithm receives loss $\ell_A^t = \ell_i^t$ corresponding to the action i that it chose at time t

Regret

Our new goal is to minimize regret.

Definition (Regret)

For all $t \in [T]$, let ℓ_A^t be the loss suffered by algorithm A at time t. Then the *regret* of algorithm A is

$$\text{Regret}(A) = \frac{1}{T}\sum_{t=1}^{T} \ell_A^t - \frac{1}{T}\min_{i \in [N]} \sum_{t=1}^{T} \ell_i^t$$

Regret

Our new goal is to minimize regret.

Definition (Regret)

For all $t \in [T]$, let ℓ_A^t be the loss suffered by algorithm A at time t. Then the *regret* of algorithm A is

$$\mathrm{Regret}(A) = \frac{1}{T}\sum_{t=1}^{T} \ell_A^t - \frac{1}{T}\min_{i \in [N]} \sum_{t=1}^{T} \ell_i^t$$

An algorithm is a *no-regret* algorithm if its regret goes to 0 as $T \to \infty$.

Multiplicative Weights Algorithm

Algorithm Multiplicative Weights (MW)

For $i \in [N]$, let $w_i^1 = 1$ be the weight of action i at time 1.
for $t = 1, \ldots, T$ **do**
  Let $W^t = \sum_{i \in [N]} w_i^t$ be the total weight at time t.
  Choose action $i \in [N]$ at random according to the distribution $D(i) = \frac{w_i^t}{W^t}$
  Pay loss ℓ_i^t for action i at time t
  Update weights: $w_j^{t+1} \leftarrow w_j^t \cdot e^{-\varepsilon \ell_j^t}$ for all $j \in [N]$
end for

$$w_i^{t+1} \leftarrow w_i^t \cdot (1 - \varepsilon \ell_i^t)$$

Multiplicative Weights Analysis

Theorem

MW is a no-regret algorithm. Specifically, it has expected regret $O(\varepsilon + \frac{\log(N)}{\varepsilon T})$. *That is,*

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\ell_A^t] - \frac{1}{T}\min_{i\in[N]}\sum_{t=1}^{T}\ell_i^t \in O\left(\varepsilon + \frac{\ln(N)}{\varepsilon T}\right)$$

Multiplicative Weights Analysis

Theorem

MW is a no-regret algorithm. Specifically, it has expected regret $O(\varepsilon + \frac{\log(N)}{\varepsilon T})$. *That is,*

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\ell_A^t] - \frac{1}{T}\min_{i\in[N]}\sum_{t=1}^{T}\ell_i^t \in O\left(\varepsilon + \frac{\ln(N)}{\varepsilon T}\right)$$

If we set $\varepsilon = \sqrt{\frac{\ln N}{T}}$, we get that the expected regret of MW is at most $2\sqrt{\frac{\ln N}{T}}$.

Multiplicative Weights Analysis

Theorem

MW is a no-regret algorithm. Specifically, it has expected regret $O(\varepsilon + \frac{\log(N)}{\varepsilon T})$. *That is,*

$$\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}[\ell_A^t] - \frac{1}{T}\min_{i\in[N]}\sum_{t=1}^{T}\ell_i^t \in O\left(\varepsilon + \frac{\ln(N)}{\varepsilon T}\right)$$

If we set $\varepsilon = \sqrt{\frac{\ln N}{T}}$, we get that the expected regret of MW is at most $2\sqrt{\frac{\ln N}{T}}$. This means that MW is a no-regret algorithm, since its regret $\to 0$ as $T \to \infty$.

Proof Sketch

Let $\boldsymbol{W_t = \sum_{i \in [N]} w_i^t}$

- We can show that $\boldsymbol{W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])}$

Proof Sketch

Let $W_t = \sum_{i \in [N]} w_i^t$

- We can show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])$

For any action i, we know that

- $W_{T+1} \geq w_i^{T+1} = \exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t)$

Proof Sketch

Let $W_t = \sum_{i \in [N]} w_i^t$

- We can show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])$

For any action i, we know that

- $W_{T+1} \geq w_i^{T+1} = \exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t)$

So, like in the analysis for Weighted Majority, we have upper and lower bounds on W_{T+1}:

$$\exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t) = w_i^{T+1} \leq W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$$

Proof Sketch

Let $\boldsymbol{W_t} = \sum_{i \in [N]} \boldsymbol{w_i^t}$

- We can show that $\boldsymbol{W_{T+1}} \leq \boldsymbol{N} \exp(\varepsilon^2 \boldsymbol{T} - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])$

For any action $\boldsymbol{i}$, we know that

- $\boldsymbol{W_{T+1}} \geq \boldsymbol{w_i^{T+1}} = \exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t)$

So, like in the analysis for Weighted Majority, we have upper and lower bounds on $\boldsymbol{W_{T+1}}$:

$$\exp\left(-\varepsilon \sum_{t=1}^{T} \ell_i^t\right) = \boldsymbol{w_i^{T+1}} \leq \boldsymbol{W_{T+1}} \leq \boldsymbol{N} \exp\left(\varepsilon^2 \boldsymbol{T} - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t]\right)$$

$$\Rightarrow -\varepsilon \sum_{t \in [T]} \ell_i^t \leq \ln \boldsymbol{N} + \varepsilon^2 \boldsymbol{T} - \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t]$$

Proof Sketch

Let $W_t = \sum_{i \in [N]} w_i^t$

- We can show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^T \mathbb{E}[\ell_A^t])$

For any action i, we know that

- $W_{T+1} \ge w_i^{T+1} = \exp(-\varepsilon \sum_{t=1}^T \ell_i^t)$

So, like in the analysis for Weighted Majority, we have upper and lower bounds on W_{T+1}:

$$\exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t) = w_i^{T+1} \le W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$$

$$\Rightarrow -\varepsilon \sum_{t \in [T]} \ell_i^t \le \ln N + \varepsilon^2 T - \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t]$$

$$\Rightarrow \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \varepsilon \sum_{t \in [T]} \ell_i^t \le \ln N + \varepsilon^2 T$$

Proof Sketch

Let $W_t = \sum_{i \in [N]} w_i^t$

- We can show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])$

For any action i, we know that

- $W_{T+1} \geq w_i^{T+1} = \exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t)$

So, like in the analysis for Weighted Majority, we have upper and lower bounds on W_{T+1}:

$$\exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t) = w_i^{T+1} \leq W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$$

$$\Rightarrow -\varepsilon \sum_{t \in [T]} \ell_i^t \leq \ln N + \varepsilon^2 T - \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t]$$

$$\Rightarrow \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \varepsilon \sum_{t \in [T]} \ell_i^t \leq \ln N + \varepsilon^2 T$$

$$\Rightarrow \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \sum_{t \in [T]} \ell_i^t \leq \frac{\ln N}{\varepsilon} + \varepsilon T$$

Proof Sketch

Let $W_t = \sum_{i \in [N]} w_i^t$

- We can show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t=1}^{T} \mathbb{E}[\ell_A^t])$

For any action i, we know that

- $W_{T+1} \geq w_i^{T+1} = \exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t)$

So, like in the analysis for Weighted Majority, we have upper and lower bounds on W_{T+1}:

$$\exp(-\varepsilon \sum_{t=1}^{T} \ell_i^t) = w_i^{T+1} \leq W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$$

$$\Rightarrow -\varepsilon \sum_{t \in [T]} \ell_i^t \leq \ln N + \varepsilon^2 T - \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t]$$

$$\Rightarrow \varepsilon \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \varepsilon \sum_{t \in [T]} \ell_i^t \leq \ln N + \varepsilon^2 T$$

$$\Rightarrow \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \sum_{t \in [T]} \ell_i^t \leq \frac{\ln N}{\varepsilon} + \varepsilon T \Rightarrow \frac{1}{T} \sum_{t \in [T]} \mathbb{E}[\ell_A^t] - \frac{1}{T} \sum_{t \in [T]} \ell_i^t \leq \frac{\ln N}{\varepsilon T} + \varepsilon$$

Proof Sketch

Still need to show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

Proof Sketch

Still need to show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t)$$

By definition of w_i^{t+1}

Proof Sketch

Still need to show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t) \qquad \text{By definition of } w_i^{t+1}$$

$$\leq \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \qquad \exp(-x) \leq 1 - x + x^2 \text{ for } x > 0$$

Proof Sketch

Still need to show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t) \qquad \text{By definition of } w_i^{t+1}$$

$$\leq \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \qquad \exp(-x) \leq 1 - x + x^2 \text{ for } x > 0$$

$$\leq \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \qquad \text{All losses in } [0, 1]$$

Proof Sketch

Still need to show that $W_{T+1} \leq N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t) \qquad \text{By definition of } w_i^{t+1}$$

$$\leq \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \qquad \exp(-x) \leq 1 - x + x^2 \text{ for } x > 0$$

$$\leq \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \qquad \text{All losses in } [0, 1]$$

$$= (1 + \varepsilon^2) \sum_{i \in [N]} w_i^t - \varepsilon \sum_{i \in [N]} w_i^t \ell_i^t$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t\in[T]}^{T} \mathbb{E}[\ell_A^t])$

$$\begin{aligned}
W_{t+1} &= \sum_{i\in[N]} w_i^{t+1} = \sum_{i\in[N]} w_i^t \exp(-\varepsilon \ell_i^t) \\
&\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \\
&\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \\
&= (1+\varepsilon^2) \sum_{i\in[N]} w_i^t - \varepsilon \sum_{i\in[N]} w_i^t \ell_i^t
\end{aligned}$$

$$= (1+\varepsilon^2) W_t - \varepsilon W_t \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t \qquad \text{by def. of } W_t$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t\in[T]}^{T} \mathbb{E}[\ell_A^t])$

$$\begin{aligned} W_{t+1} &= \sum_{i\in[N]} w_i^{t+1} = \sum_{i\in[N]} w_i^t \exp(-\varepsilon \ell_i^t) \\ &\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \\ &\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \\ &= (1+\varepsilon^2) \sum_{i\in[N]} w_i^t - \varepsilon \sum_{i\in[N]} w_i^t \ell_i^t \end{aligned}$$

$$\begin{aligned} &= (1+\varepsilon^2) W_t - \varepsilon W_t \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t && \text{by def. of } W_t \\ &= W_t (1 + \varepsilon^2 - \varepsilon \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t) \end{aligned}$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t)$$

$$\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2)$$

$$\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2)$$

$$= (1 + \varepsilon^2) \sum_{i \in [N]} w_i^t - \varepsilon \sum_{i \in [N]} w_i^t \ell_i^t$$

$$= (1 + \varepsilon^2) W_t - \varepsilon W_t \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t \qquad \text{by def. of } W_t$$

$$= W_t (1 + \varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t)$$

$$\le W_t \exp(\varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t) \qquad \exp(x) \ge 1 + x$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$
\begin{aligned}
W_{t+1} &= \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t) \\
&\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \\
&\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \\
&= (1 + \varepsilon^2) \sum_{i \in [N]} w_i^t - \varepsilon \sum_{i \in [N]} w_i^t \ell_i^t
\end{aligned}
$$

$$
\begin{aligned}
&= (1 + \varepsilon^2) W_t - \varepsilon W_t \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t && \text{by def. of } W_t \\
&= W_t (1 + \varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t) \\
&\le W_t \exp(\varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t) && \exp(x) \ge 1 + x \\
&= W_t \exp(\varepsilon^2 - \varepsilon \mathbb{E}[\ell_A^t]) && \text{by def. of } \mathbb{E}[\ell_A^t]
\end{aligned}
$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t \in [T]}^{T} \mathbb{E}[\ell_A^t])$

$$
\begin{aligned}
W_{t+1} &= \sum_{i \in [N]} w_i^{t+1} = \sum_{i \in [N]} w_i^t \exp(-\varepsilon \ell_i^t) \\
&\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2) \\
&\le \sum_{i \in [N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2) \\
&= (1 + \varepsilon^2) \sum_{i \in [N]} w_i^t - \varepsilon \sum_{i \in [N]} w_i^t \ell_i^t
\end{aligned}
$$

$$
\begin{aligned}
&= (1 + \varepsilon^2) W_t - \varepsilon W_t \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t && \text{by def. of } W_t \\
&= W_t (1 + \varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t) \\
&\le W_t \exp(\varepsilon^2 - \varepsilon \sum_{i \in [N]} \frac{w_i^t}{W_t} \ell_i^t) && \exp(x) \ge 1 + x \\
&= W_t \exp(\varepsilon^2 - \varepsilon \mathbb{E}[\ell_A^t]) && \text{by def. of } \mathbb{E}[\ell_A^t]
\end{aligned}
$$

Proof Sketch

Still need to show that $W_{T+1} \le N \exp(\varepsilon^2 T - \varepsilon \sum_{t\in[T]}^{T} \mathbb{E}[\ell_A^t])$

$$W_{t+1} = \sum_{i\in[N]} w_i^{t+1} = \sum_{i\in[N]} w_i^t \exp(-\varepsilon \ell_i^t)$$

$$\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2 {\ell_i^t}^2)$$

$$\le \sum_{i\in[N]} w_i^t \cdot (1 - \varepsilon \ell_i^t + \varepsilon^2)$$

$$= (1+\varepsilon^2) \sum_{i\in[N]} w_i^t - \varepsilon \sum_{i\in[N]} w_i^t \ell_i^t$$

$$= (1+\varepsilon^2) W_t - \varepsilon W_t \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t \qquad \text{by def. of } W_t$$

$$= W_t (1 + \varepsilon^2 - \varepsilon \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t)$$

$$\le W_t \exp(\varepsilon^2 - \varepsilon \sum_{i\in[N]} \frac{w_i^t}{W_t} \ell_i^t) \qquad \exp(x) \ge 1 + x$$

$$= W_t \exp(\varepsilon^2 - \varepsilon \mathbb{E}[\ell_A^t]) \qquad \text{by def. of } \mathbb{E}[\ell_A^t]$$

Unrolling over all $t \in [T]$:

$$\Rightarrow W_{T+1} \le W_1 \cdot \Pi_{t\in[T]} \exp(\varepsilon^2 - \varepsilon \mathbb{E}[\ell_A^t])$$
$$= N \exp(\varepsilon^2 T - \varepsilon \sum_{t\in[T]} \mathbb{E}[\ell_A^t])$$

Summary

- Online learning models learning problems where data arrives "online" (one at a time)
- Given a set of actions to choose from, we want to learn a good sequence of actions to take so that we do not incur too much loss
- We can analyze the performance of online learning algorithms using the notion of *regret*: how well did the algorithm perform compared to the best action in hindsight?
- We showed the multiplicative weights algorithm is a no-regret algorithm (its expected regret goes to $\mathbf{0}$ as $T \to \infty$).