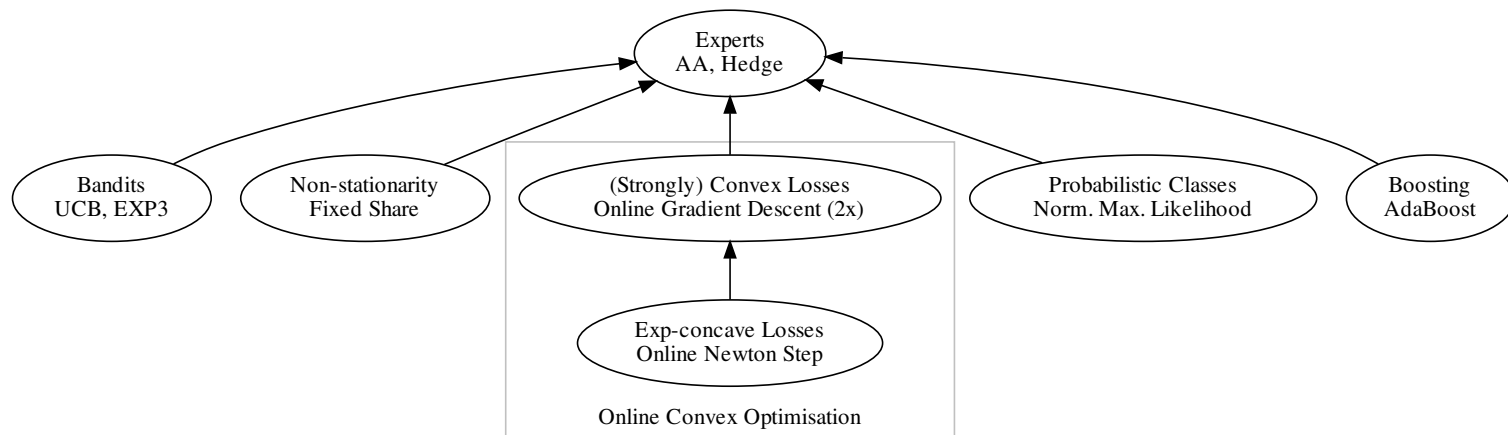


Machine Learning Theory. Lecture 8.

Wouter M. Koolen

- Online Learning, Experts, basic algorithms



Online learning focus

- Tight feedback loop (recurring prediction task)
- Continuous learning (no training/learning separation)
- Adversarial analysis (Prequential principle, individual sequence. There is only the data. Also establishes robustness of statistical estimators.)
- Emphasis on both computational and statistical performance
- Regret: relative notion of performance

Application domains

Truly sequential problems:

- electricity demand prediction (EDF, also Amazon)
- mobile device power management
- hybrid cars engine switching
- caching
- medical trials (bandits)
- online advertisement (bandits)
- weather forecasting
- data compression (CTW)
- statistical testing
- investment (Universal portfolios)
- input assistants (e.g. Dasher)
- prediction with expert advice (meld human and machine prediction)
- online convex optimisation

Wider application

- Big data sets (transport state of online algorithm instead of data, online to batch conversion)
- Convex optimisation
- Game theory (online learning methods for approximate equilibrium)
- General understanding
 - Uncertainty and ways to manipulate it
 - Makeup of and patterns in data
 - Complexity of classes of strategies

The menu for today

Two fundamental and prototypical online learning problems

- The mix-loss game
 - Aggregating Algorithm
- The dot-loss game
 - Hedge Algorithm

Mix-loss game

Protocol:

- For $t = 1, 2, \dots$
 - Learner chooses a distribution $w_t \in \Delta_K$ on K “experts”.
 - Adversary reveals loss vector $\ell_t \in (-\infty, \infty]^K$.
 - Learner’s loss is the **mix loss** $-\ln \left(\sum_{k=1}^K w_t^k e^{-\ell_t^k} \right)$

Instances:

- Investment (loss is *negative log-growth*)
- Data compression (loss is *code length*)
- Probability forecasting (loss is *logarithmic loss*)

Mix-loss objective

Obviously we cannot guarantee small loss.

Idea: relative evaluation, i.e. seek performance close to best expert.

Definition: After T rounds of the mix-loss game, the *regret* is given by

$$R_T = \underbrace{\sum_{t=1}^T -\ln \left(\sum_{k=1}^K w_t^k e^{-\ell_t^k} \right)}_{\text{Learner's mix loss}} - \underbrace{\min_k \sum_{t=1}^T \ell_t^k}_{\text{loss of best expert}}$$

Goal: design a strategy for Learner that guarantees low regret.

Worst-case regret and Minimax regret

An strategy for the learner assigns a next action w_t to each history $(w_1, \ell_1), \dots, (w_{t-1}, \ell_{t-1})$.

Definition: The *worst-case regret* of a strategy S for the learner is

$$\max_{\ell_1}, \dots, \max_{\ell_T} R_T$$

where the w_t are chosen according to S .

Definition: The *minimax regret* of the mix loss game is

$$\min_{w_1} \max_{\ell_1}, \dots, \min_{w_T} \max_{\ell_T} R_T$$

Mix-loss regret: lower bound (adversary construction)

Theorem: Any strategy for Learner has worst-case regret $\geq \ln K$, already in $T = 1$ round.

Idea: Look at $k_{\text{low}} \in \arg \min_k w_1^k$ so that $w_1^{k_{\text{low}}} \leq \frac{1}{K}$.

Administer loss killing everyone but k_{low}

$$\ell_1^k = \begin{cases} \infty & k \neq k_{\text{low}} \\ 0 & k = k_{\text{low}} \end{cases}$$

Now Learner's mix loss equals

$$-\ln \left(\sum_{k=1}^K w_1^k e^{-\ell_1^k} \right) = -\ln \left(w_1^{k_{\text{low}}} e^{-\ell_1^{k_{\text{low}}}} \right) \geq \ln K + \ell_1^{k_{\text{low}}}$$

The Aggregating Algorithm for mix loss

Definition: The *Aggregating Algorithm* plays weights in round t :

$$w_t^k = \frac{e^{-\sum_{s=1}^{t-1} \ell_s^k}}{\sum_{j=1}^K e^{-\sum_{s=1}^{t-1} \ell_s^j}} \quad (\text{AA})$$

or, equivalently, $w_1^k = \frac{1}{K}$ and

$$w_{t+1}^k = \frac{w_t^k e^{-\ell_t^k}}{\sum_{j=1}^K w_t^j e^{-\ell_t^j}} \quad (\text{AA, incremental})$$

Many names

- (Generalisation of) Bayes' rule
- Exponentially weighted average

Mix-loss regret: upper bound (algorithm)

Theorem: The regret of the Aggregating Algorithm does not exceed $R_T \leq \ln K$ for all $T \geq 0$.

Proof: Crucial observation is that mix loss *telescopes*

$$\begin{aligned} \sum_{t=1}^T -\ln \left(\sum_{k=1}^K w_t^k e^{-\ell_t^k} \right) &= \sum_{t=1}^T -\ln \left(\sum_{k=1}^K \frac{e^{-\sum_{s=1}^{t-1} \ell_s^k}}{\sum_{j=1}^K e^{-\sum_{s=1}^{t-1} \ell_s^j}} e^{-\ell_t^k} \right) \\ &= \sum_{t=1}^T -\ln \left(\frac{\sum_{k=1}^K e^{-\sum_{s=1}^t \ell_s^k}}{\sum_{j=1}^K e^{-\sum_{s=1}^{t-1} \ell_s^j}} \right) \\ &= -\ln \left(\sum_{k=1}^K e^{-\sum_{t=1}^T \ell_t^k} \right) + \ln K. \end{aligned}$$

Bounding the sum from below by a max results in

$$\leq \min_k \sum_{t=1}^T \ell_t^k + \ln K \quad (1)$$

Dot-loss game

Protocol:

- For $t = 1, 2, \dots$
 - Learner chooses a distribution $w_t \in \Delta_K$ on K “experts”.
 - Adversary reveals loss vector $\ell_t \in [0, 1]^K$.
 - Learner’s loss is the **dot loss** $w_t^\top \ell_t = \sum_{k=1}^K w_t^k \ell_t^k$

Many names:

- Decision Theoretic Online Learning
- Prediction with Expert Advice
- The Hedge setting

Dot-loss objective

Definition: *Regret* after T rounds:

$$R_T = \sum_{t=1}^T \mathbf{w}_t^\top \ell_t - \min_k \sum_{t=1}^T \ell_t^k$$

Goal: design an algorithm for Learner that guarantees low regret.

Mix loss vs Dot-loss (Jensen)

By Jensen's Inequality for the convex function $x \mapsto -\ln(x)$

$$\underbrace{-\ln\left(\sum_{k=1}^K w_t^k e^{-\ell_t^k}\right)}_{\text{mix loss}} \leq \underbrace{\sum_{k=1}^K w_t^k \ell_t^k}_{\text{dot loss}} \quad (2)$$

So the dot loss game is harsher for the Learner ...

...but maybe we can find a converse inequality (with small overhead)

Mix loss vs Dot loss (Hoeffding)

Lemma 1 (Hoeffding). Fix zero-mean r.v. $X \in [a, b]$, and let $\eta \in \mathbb{R}$. Then

$$\mathbb{E}[e^{\eta X}] \leq e^{\eta^2(b-a)^2/8}$$

(Note: Lemma is main ingredient of but not equal to Hoeffding's Bound)

Application: Fix $\mathbf{w}_t \in \Delta_K$ and $\ell_t \in [0, 1]^K$. Define r.v. X to take value $\mathbf{w}_t^\top \ell_t - \ell_t^k$ with probability w_t^k for all $k = 1, \dots, K$. Then X has mean zero, and takes values in an interval of length 1. So

$$\sum_k w_t^k e^{\eta(\mathbf{w}_t^\top \ell_t - \ell_t^k)} \leq e^{\eta^2/8}$$

and hence we obtain a converse to (2):

$$\underbrace{\mathbf{w}_t^\top \ell_t}_{\text{dot loss}} \leq \underbrace{-\frac{1}{\eta} \ln \left(\sum_k w_t^k e^{-\eta \ell_t^k} \right)}_{\eta\text{-scaled mix loss}} + \frac{\eta}{8}$$

Hedge algorithm

Idea: re-use AA for mix loss, now with *learning rate* $\eta > 0$.

Definition: The *Hedge algorithm* with *learning rate* η plays weights in round t :

$$w_t^k = \frac{e^{-\eta \sum_{s=1}^{t-1} \ell_s^k}}{\sum_{j=1}^K e^{-\eta \sum_{s=1}^{t-1} \ell_s^j}}. \quad (\text{Hedge})$$

or, equivalently, $w_1^k = \frac{1}{K}$ and

$$w_{t+1}^k = \frac{w_t^k e^{-\eta \ell_t^k}}{\sum_{j=1}^K w_t^j e^{-\eta \ell_t^j}} \quad (\text{Hedge, incremental})$$

Hedge analysis

Lemma: The regret of Hedge is bounded by

$$R_T \leq \frac{\ln K}{\eta} + T \frac{\eta}{8}$$

Proof: Applying Hoeffding's Lemma to the loss of each round gives

$$\sum_{t=1}^T \mathbf{w}_t^\top \ell_t \leq \underbrace{\sum_{t=1}^T \left(\frac{-1}{\eta} \ln \left(\sum_{k=1}^K w_t^k e^{-\eta \ell_t^k} \right) \right)}_{\text{mix loss}} + \underbrace{\eta/8}_{\text{overhead}}$$

The mix loss telescopes, and is bounded by (1) by

$$\sum_{t=1}^T \frac{-1}{\eta} \ln \left(\sum_{k=1}^K w_t^k e^{-\eta \ell_t^k} \right) \leq \min_k \sum_{t=1}^T \ell_t^k + \frac{\ln K}{\eta}. \quad (3)$$

Hedge tuning

Theorem: The Hedge regret bound is minimised at $\eta = \sqrt{\frac{8 \ln K}{T}}$ where it states

$$R_T \leq \sqrt{T/2 \ln K}.$$

Proof: pick η to cancel the derivative.

Note: tuning requires knowledge of the time horizon T .