Machine Learning Theory. Lecture 10.

Wouter M. Koolen

Recap:

- Mix loss (now with non-uniform regret bounds)

Non-stationary environments

- Switching (Fixed Share algorithm)

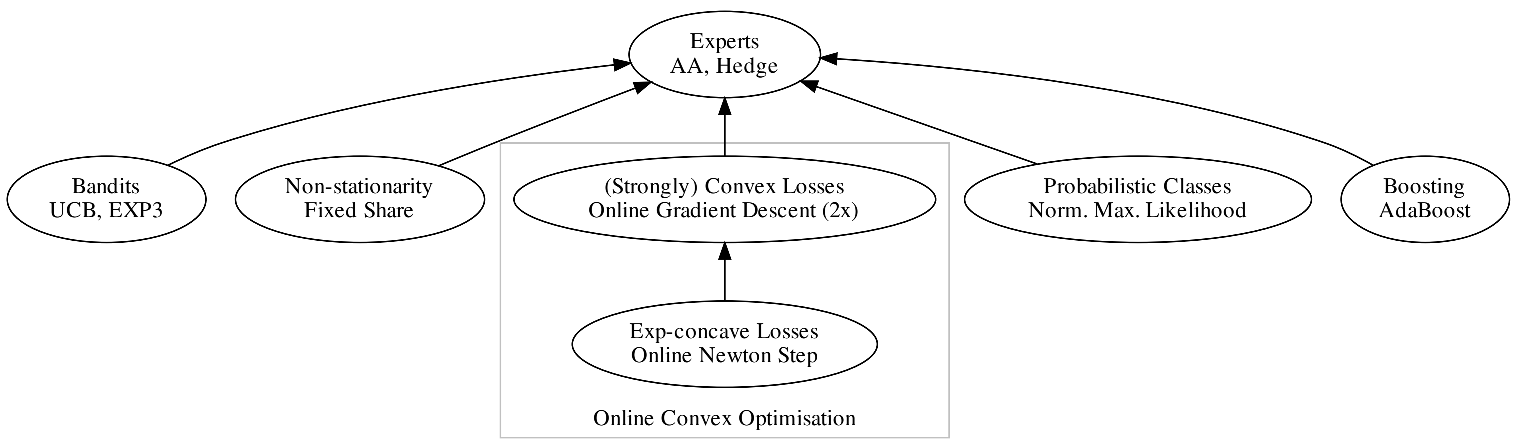

Mix loss prediction (w. non-uniform regret objective)

For $t = 1, 2, \ldots$

1. Play $\boldsymbol{w}_t \in \triangle_K$.
2. See $\boldsymbol{\ell}_t \in \mathbb{R}^K$.
3. Incur *mix loss* $\hat{\ell}_t := -\ln\left(\sum_{k=1}^K w_t^k e^{-\ell_t^k}\right)$.

Definition 1. *The* regret *w.r.t. expert $k \in [K]$ after $T \geq 0$ rounds is*

$$R_T^k := \sum_{t=1}^T \left(\hat{\ell}_t - \ell_t^k\right).$$

The Aggregating Algorithm with prior

Definition 2. *The Aggregating Algorithm with* prior $\boldsymbol{\pi} \in \triangle_K$ *plays*

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

(so $\boldsymbol{w}_1 = \boldsymbol{\pi}$ *and* $w_{t+1}^k = \frac{w_t^k e^{-\ell_t^k}}{\sum_{j=1}^K w_t^j e^{-\ell_t^j}}$*)*

Theorem 3. *The regret of AA*$-\boldsymbol{\pi}$ *w.r.t. expert* $k \in [K]$ *satisfies*

$$R_T^k \leq -\ln \pi^k$$

Proof: part of Homework 10.

Non-stationary data

1. Switching, Tracking, Shifting

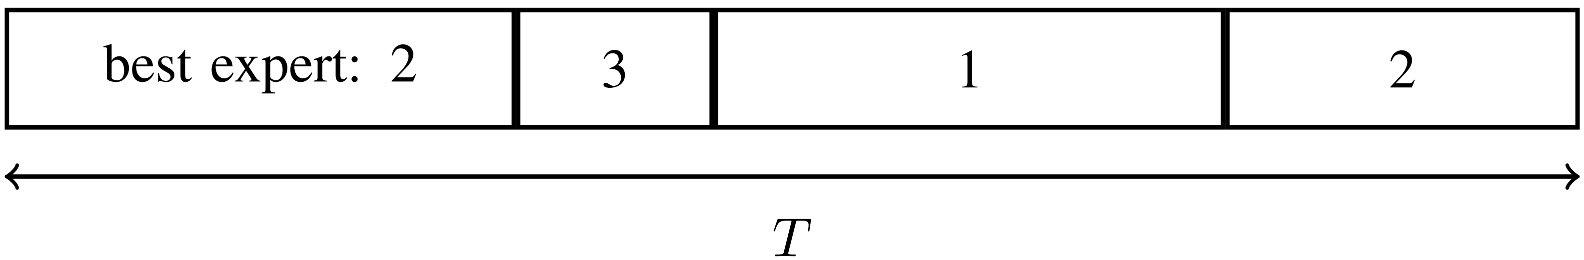

What if no single expert is good for all T rounds, but data can be split into 4 blocks with different best experts?

Non-stationary data

So far we have been looking at regret compared to a *fixed* expert/action.

$$R_T \;=\; \max_{k \in [K]} \sum_{t=1}^{T} \left(\hat{\ell}_t - \ell_t^k \right).$$

But what if we do not expect a single expert to be good for all data?

Definition 4. *Let $\xi \in [K]^T$ be a sequence of experts. The sequence regret w.r.t. ξ is defined by*

$$R_\xi \;:=\; \sum_{t=1}^{T} \left(\hat{\ell}_t - \ell_t^{\xi_t} \right)$$

Question: Can we keep the sequence regret small w.r.t. *every sequence*? Or, at least, w.r.t. many interesting sequences?

Fixed Share (defined by reduction)

Starting with K experts, create an "explosion" $[K]^T$ of *expert sequences*.

Fix *switching rate* $\alpha \in [0,1]$. Run the AA with prior $\boldsymbol{\pi} \in \triangle_{[K]^T}$ and losses $\boldsymbol{\ell}_t \in (-\infty, \infty]^{[K]^T}$ defined by

$$\pi^\xi := \frac{1}{K} \prod_{t=2}^{T} \begin{cases} 1-\alpha & \text{if } \xi_{t-1} = \xi_t \\ \frac{\alpha}{K-1} & \text{if } \xi_{t-1} \neq \xi_t \end{cases}$$

$$\ell_t^\xi := \ell_t^{\xi_t}$$

Based on the AA predictions $w_t^\xi \in \triangle_{[K]^T}$, *Fixed Share* plays

$$w_t^k := \sum_{\xi \in [K]^T : \xi_t = k} w_t^\xi \qquad \text{(FS)}$$

Crucial equality:

$$\underbrace{-\ln\sum_{\xi} w_t^{\xi} e^{-\ell_t^{\xi}}}_{\text{mix loss of AA}} \;=\; \underbrace{-\ln\sum_{k} w_t^{k} e^{-\ell_t^{k}}}_{\text{mix loss of FS}}$$

So we can directly apply AA regret bound to obtain FS regret bound.

Fixed Share: Regret Bound

Application of the AA regret bound, (Theorem 3) gives

Theorem 5. *Fixed Share ensures that the regret on each sequence ξ with $B \leq T$ contiguous blocks is at most*

$$
\begin{aligned}
R_\xi &\leq -\ln \pi^\xi \\
&= \underbrace{\ln K + (B-1)\ln(K-1)}_{\textit{expert labelling cost}} \underbrace{-(B-1)\ln\alpha - (T-B)\ln(1-\alpha)}_{\textit{switching location cost}}.
\end{aligned}
$$

Corollary 6. *If we know the number of blocks B in advance, we can optimise the bound by setting $\alpha = \frac{B-1}{T-1}$ to find*

$$
R_\xi \leq \ln K + (B-1)\ln(K-1) + (T-1)H\left(\frac{B-1}{T-1}\right),
$$

where $H(\alpha) = -\alpha\ln\alpha - (1-\alpha)\ln(1-\alpha)$ is the binary entropy.

Fixed Share: Computation Collapses

Seems we need to maintain exponentially many weights. But prior π^ξ is *Markov*

Theorem 7. *The weights of Fixed Share with switching rate α can be computed incrementally in $\mathcal{O}(K)$ time per round (same as AA) as*

$$w_{t+1}^k \;=\; (1-\beta)\frac{w_t^k e^{-\ell_t^k}}{\sum_j w_t^j e^{-\ell_t^j}} + \frac{\beta}{K}$$

where $\beta = \alpha \frac{K}{K-1}$.

Proof: part of Homework 10.

We see that the Fixed Share update is a weighted combination of the incremental AA update and the uniform prior.

Conclusion

Technique for adapting to changing environment

- Fixed Share for switching between experts

Conceptual message:

- Adapting to changing environment is not automatic
- Modelling with explicit sequences