The Hedge Algorithm on a Continuum

Walid Krichene   Maximilian Balandat   Claire Tomlin   Alexandre Bayen

Electrical Engineering and Computer Sciences, UC Berkeley

{krichene, balandat, tomlin, bayen}@eecs.berkeley.edu

International Conference on Machine Learning

Online learning on a continuum

Online decision problem on S.

1: **for** $t \in \mathbb{N}$ **do**
2: Decision maker chooses distribution $x^{(t)}$ over S.
3: A loss function $\ell^{(t)} : S \to [0, M]$ is revealed. Assumed L-Lipschitz.
4: The decision maker incurs expected loss

$$\left\langle x^{(t)}, \ell^{(t)} \right\rangle = \int_S x^{(t)}(s)\ell^{(t)}(s)\lambda(ds) = \mathbb{E}_{s\sim x^{(t)}}[\ell^{(t)}(s)]$$

Regret

$$R^{(T)}(x) = \sum_{t=1}^{T} \left\langle x^{(t)}, \ell^{(t)} \right\rangle - \left\langle x, \sum_{t=1}^{T} \ell^{(t)} \right\rangle$$

Objective: design algorithm with

$$\sup_{(\ell^{(t)})} \sup_{x\in\Delta^N} R^{(T)}(x) = o(T)$$

Regret rates

Variant of this problem: Online optimization on convex sets.

Assumptions on $\ell^{(t)}$	convex	α-exp-concave	uniformly L-Lipschitz
Assumptions on S	convex	convex	v-uniformly fat
Method	Gradient (Zinkevich)	Hedge (Hazan et al.)	Hedge (this work)
Learning rates	$1/\sqrt{t}$	α	$1/\sqrt{t}$
$R^{(t)}$	$\mathcal{O}(\sqrt{t})$	$\mathcal{O}(\log t)$	$\mathcal{O}(\sqrt{t \log t})$

Table: Regret upper bounds for different classes of losses.

Hedge on a continuum

Online decision problem over a compact set S.

1: **for** $t \in \mathbb{N}$ **do**
2: Play $\sim x^{(t)}$
3: Reveal $\ell^{(t)} : S \to [0, M]$. Call $L^{(t)} = \sum_{\tau=1}^{t} \ell^{(\tau)}$
4: Update

$$x^{(t+1)}(s) \propto x^{(0)}(s)e^{-\eta_{t+1}L^{(t)}(s)}$$

Interpretation: instance of the dual averaging method

$$x^{(t+1)} \in \arg\min_{x\in\Delta(S)} \left\langle L^{(t)}, x \right\rangle + \frac{1}{\eta_{t+1}}\psi(x)$$

with

- Hilbert space $\mathcal{H} = L^2(S)$, $\langle \ell, x\rangle = \int_S \ell(s)x(s)\lambda(ds)$
- $\Delta(S) = \{x \in L^2(S) : x \geq 0,\ \|x\|_1 = 1\}$
- $\psi(x) = \int_S x(s) \ln x(s)\lambda(ds)$

Basic Regret Bound

Basic regret bound

For all $x \in \Delta(S)$,

$$R^{(T)}(x) \leq \frac{M^2}{2}\sum_{\tau=1}^{t}\eta_{\tau+1} + \frac{\psi(x)}{\eta_{t+1}}$$

In the finite case, $\psi(x) \leq \ln N$ on Δ^N.

But ψ is unbounded on $\Delta(S)$.

Take $x = \frac{1}{\lambda(A)}1_A$ for some $A \subset S$. Then $\psi(x) = \int_S x(s)\ln x(s)\lambda(ds) = \ln\frac{1}{\lambda(A)}$.

Working around unbounded regularizers

For any density y

$$R^{(t)}(x) \leq R^{(t)}(y) + \left\langle L^{(t)}, y - \delta_{s_t^\star}\right\rangle$$

In particular, if y is uniform on A_t which contains $s_t^\star = \arg\min_{s\in S} L^{(t)}(s)$,

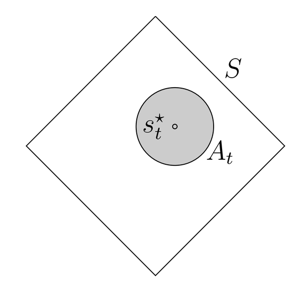

$$\sup_{x\in\Delta(S)} R^{(t)}(x) \leq R^{(t)}(y) + \left\langle L^{(t)}, y - \delta_{s_t^\star}\right\rangle$$

$$\leq \frac{M^2}{2}\sum_{\tau=1}^{t}\eta_{\tau+1} + \frac{1}{\eta_{t+1}}\ln\frac{1}{\lambda(A_t)} + Ltd(A_t)$$

Uniformly fat sets

Uniformly fat sets

S is v-uniformly fat (w.r.t. the measure λ) if $\forall s \in S$, $\exists$ convex $K_s \subset S$, with $s \in K_s$ and $\lambda(K_s) \geq v$.

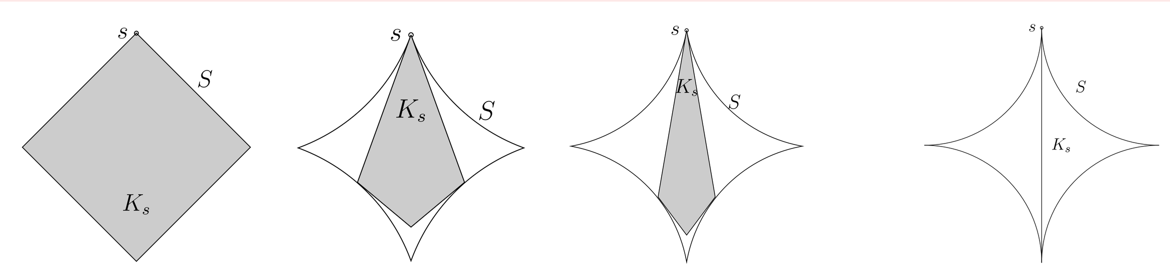

Figure: Examples of uniformly fat sets (left) and a not uniformly fat set (right)

Regret bound on uniformly fat sets

Take $A_t = s_t^\star + d_t(K_{s_t^\star} - s_t^\star)$. Then

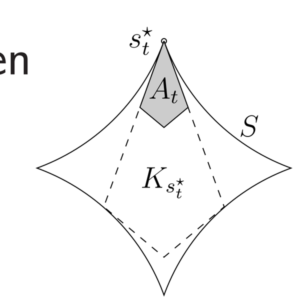

- $\lambda(A_t) \geq d_t^n v$
- $d(A_t) \leq d_t d(S)$.

Final bound

$$\sup_{x\in\Delta(S)} R^{(t)}(x) \leq \frac{M^2}{2}\sum_{\tau=1}^{t}\eta_{\tau+1} + \frac{\ln\frac{1}{v}}{\eta_{t+1}} + \frac{n\ln t}{\eta_{t+1}} + Ld(S)$$

Can optimize over η_t to get

$$\sup_{x\in\Delta(S)} R^{(t)}(x) \leq Ld(S) + M\sqrt{\frac{n\ln t + \ln\frac{1}{v}}{2}}\sqrt{t}$$

Generalizing to dual averaging

Dual averaging with learning rates (η_t), strongly convex regularizer ψ

1: **for** $t \in \mathbb{N}$ **do**
2: Play $x^{(t)}$
3: Discover $\ell^{(t)} \in \mathcal{H}^*$
4: Update $x^{(t+1)} = \arg\min_{x\in\Delta(S)} \left\langle L^{(t)}, x\right\rangle + \frac{1}{\eta_{t+1}}$.

- $\mathcal{H}$ is infinite dimensional. Can we solve $\min_{x\in\Delta(S)} \left\langle L^{(t)}, x\right\rangle + \frac{1}{\eta_{t+1}}\psi(x)$?
- Can we obtain a sublinear regret bound?

Yes, for a family of strongly convex f-divergences.

Numerical experiments

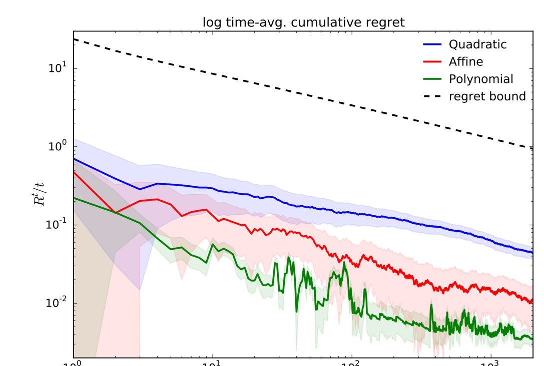

Figure: Per-round regret

Hedge algorithm
- on set S
- with Lipschitz losses
- with $\eta_t = \Theta\sqrt{\frac{\ln t}{t}}$

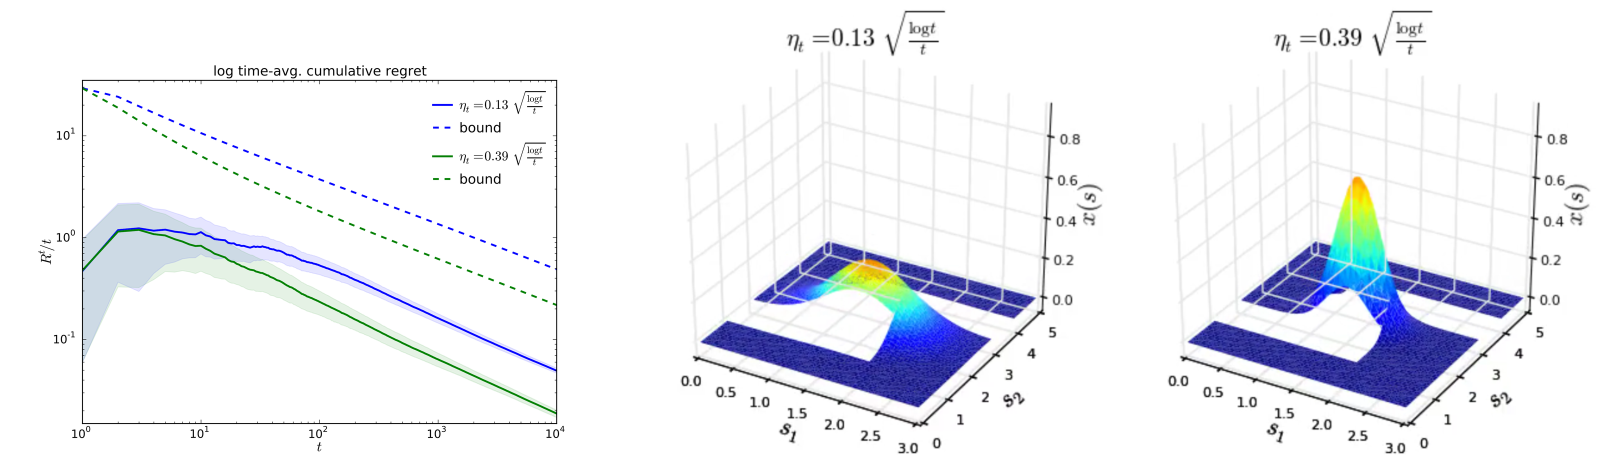

Figure: Effect of the learning rate $\eta_t = \theta\sqrt{\frac{\ln t}{t}}$

Learning on a cover

- Given a horizon T and a cover $\mathcal{A}_T$ with $d(A) \leq d_T d(S)$ for all $A \in \mathcal{A}_T$.
- Run discrete Hedge on elements of the cover.
- Then

$$R^{(T)}(x) \leq \underbrace{\frac{M^2T\eta}{8} + \frac{\ln|\mathcal{A}_T|}{\eta}}_{\text{Discrete Hedge}} + \underbrace{Ld(S)d_T}_{\text{Additional regret}}$$

- With $|\mathcal{A}_T| \approx \frac{1}{d_T^n}$, $R^{(T)}(x) \leq \frac{M^2T\eta}{8} + \frac{\ln\frac{1}{d_T^n}}{\eta} + LD(S)d_T$.
- Have to explicitly compute a (hierarchical) cover.

Conclusion

Summary

- Can learn on a continuum, when losses are Lipschitz and S has reasonable geometry.
- Similar guarantee to learning on a cover, but do not need to maintain a cover.
- Can generalize to the dual averaging method.

Extensions and open questions

- Bandit formulation.
- Regret lower bound.
- When is it easy to sample from the Hedge distribution?

This work was supported in part by FORCES (Foundations Of Resilient CybEr-physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166).