New Bounds for Geometric-Stopping Version of Prediction with Expert Advice

Vladimir A. Kobzar[1], Robert V. Kohn[2] and Zhilei Wang[2]

[1]NYU Center for Data Science, [2]Courant Institute of Mathematical Sciences

Prediction with expert advice

In each $t \in [T]$,

- the *player* determines the mix of N experts to follow - distribution $p_t \in \Delta_N$;
- the *adversary* allocates losses to them - distribution a_t over $[-1,1]^N$; and
- expert losses $q_t \sim a_t$, player's choice of expert $I_t \sim p_t$; these samples revealed to both parties.

Our contribution

Previously we developed a PDE viewpoint for the *fixed horizon* (FH) version of the problem where the *stopping time* T is fixed (COLT 2020)

This paper (MSML 2020) extends this viewpoint to the *geometric stopping* (GS) version where the stopping time $T \sim G$ and G =Geom(mean $\frac{1}{\delta}$)

- Specifically, if an FH adversary does not depend on time (stationary), it can be used for GS
- *Technically*: Given a FH potential, its Laplace transform gives a GS potential
- *Intuition*: This transform is the expectation w/r/t the Exp distribution (limit of G when $\delta \to 0$)
- *Key result:* Obtain the first lower bounds for general N for GS

Definitions

- *Instantaneous regret*: $r_\tau = q_{I_\tau,\tau}\mathbb{1} - q_\tau$
- *Accumulated regret*: $x_t = \Sigma_{\tau<t} r_\tau$
- *Final regret*: FH – $R_T(p,a) = \mathbb{E}_{p,a} \max_i x_{i,T}$; GS – $R(p,a) \ = \mathbb{E}_G R_T(p,a)$

Acknowledgements

NSF grant DMS-1311833; Moore-Sloan Data Science Environment at NYU

We use the value function

- Focus first on player strategies/upper bounds
- Assume player p is Markovian: depends only on x
- *Value function* v_p: expected final-time regret achieved by p if the game starts with realized regret x (and the adversary behaves optimally)
- Characterized by

$$v_p(x) = \delta \max_i x_i + (1-\delta) \max_a \mathbb{E}_{a,p}\, v_p(x+r)$$

Upper bound potentials/players

A function $\hat{w} : \mathbb{R}^N \to \mathbb{R}$, nondecr. in x_i, which solves

$$\begin{cases} \hat{w}(x) \geq \max_i x_i + \frac{1-\delta}{2\delta} \max_{q\in[-1,1]^N} \langle D^2\hat{w}(x)\cdot q, q\rangle \\ \hat{w}(x+c\mathbb{1}) = \hat{w}(x) + c \end{cases}$$

- The associated player $p = \nabla \hat{w}$
- Leads to an upper bound on v_p if $\hat{w}(x) - \max_i x_i$ is uniformly bounded below
- $\Rightarrow$ regret upper bound since $v_p(0) = \max_a R(a,p)$

Results

- Provide easily-checked conditions for a func. to be useful as a lower bound or an upper bound potential
- Using the Laplace transform, construct potentials for the geometric stopping problem from potentials used for the fixed horizon version
- Obtain the first known lower bound in the geometric setting for general N associated with a simple randomized strategy

FH exponential weights potential

$w^e(x,t) = \Phi(x) + kt$ where $\Phi(x) = \frac{1}{\eta}\log(\Sigma_{i=1}^N e^{\eta x_i})$

- Associated with the player $p^e = \nabla w^e$
- The standard upper bound: $\max_a R_T(a,p^e) \leq \Phi(0) + \frac{1}{2}\eta T$.
- Thus, taking $k = \frac{1}{2}\eta$ ensures that $\max_a R_T(a,p^e) \leq w^e(0,T)$ for FH

Laplace tr.: FH→GS potential

Illustrate by the exponential weights example: $\hat{w}^e(x) = \int_0^\infty e^{-t} w^e(x,t)dt = \Phi(x) + k$

- $\Phi(x) \geq \max_i x_i$ and $\langle D^2\Phi \cdot q, q\rangle \leq \eta$
- Also $\Phi(x + c\mathbb{1}) = \Phi(x) + c$
- Taking $k = \frac{1-\delta}{2\delta}\eta$ ensures $\hat{w}^e$ satisfies our def'n of a GS upper bound potential
- Since Φ is convex, $0 \leq \langle D^2\Phi\cdot q, q\rangle$. Thus $\hat{w}^e(x) - \max_i x_i \geq 0$.

Proof of $v_p \leq \hat{w}$: Idea

- *Issue*: want to use induction backwards ("verification" argument), but don't know T
- *Sol'n*: introduce a new problem, which is the same except that it ends at t_0 (if it doesn't end earlier in accordance with the GS condition)
- The difference in regret relative to the original problem→ 0 as $t_0 \to \infty$.
- Suffices to bound the value g of the new problem.
- It is given by a dynamic program: $g(x,t_0) = \max_i x_i$ an,d for $t \leq t_0 - 1$,

$$g(x,t) = \delta \max_i x_i + (1-\delta) \min_p \mathbb{E}_{a,p}\, g(x+r, t+1)$$

Proof of $v_p \leq \hat{w}$: "verification" arg.

1. Control increase of $\hat{w}$ as the game evolves: the choice $p = \nabla\hat{w}$ eliminates the 1st-order Taylor term in this evolution for all q
2. Show $g \leq \hat{w}$ by induction (and thus $v_p \leq \hat{w}$)

Lower bound potentials/adversaries

- Adversary a Markovian & "balanced": $\mathbb{E}_a q_i = \mathbb{E}_a q_j$
- Use the value function v_a for this adversary
- *Lower bound potential* is a function $\hat{u} : \mathbb{R}^N \to \mathbb{R}$ which solves

$$\begin{cases} \hat{u} \leq \max_i x_i + \frac{1-\delta}{2\delta}\mathbb{E}_a\langle D^2\hat{u}(x)\cdot q, q\rangle \\ \hat{u}(x+c\mathbb{1}) = \hat{u}(x) + c \end{cases}$$

- $\hat{u} \leq v_a$ (modulo error E from higher order terms)
- Regret bound $\hat{u}(0) - E \leq v_a(0) = \min_p R(a,p)$
- In estimating the expected value of $u(x+r) - u(x)$, the dependence on p is in the 1st-order Taylor term, which gets eliminated since a is balanced
- The dependence on a remains at the 2nd order

Heat-based adversary

- a^h is a uniform distribution over the following set S $\left\{q \in \{\pm 1\}^N \mid \Sigma_{i=1}^N q_i = \pm 1\right\}$ for N odd or $\left\{q \in \{\pm 1\}^N \mid \Sigma_{i=1}^N q_i = 0\right\}$ for N even
- Potential $\hat{u}$ is the Laplace transform of the sol'n of the linear heat equation

$$\begin{cases} u_t + \kappa\Delta u = 0 \\ u(x,0) = \max_i x \end{cases} \quad u(x,t) = \alpha \int e^{-\frac{\|y\|^2}{2\sigma^2}} \max_k (x_k - y_k) dy$$

where $\alpha = (2\pi\sigma^2)^{-\frac{N}{2}}$ and $\sigma^2 = -2\kappa t$.

- Satisfies our def'n of a lower bound potential for a well-chosen κ
- The leading order asymptotics of our lower bound $\hat{u}(0) = \Omega\left(\sqrt{\frac{\log N}{\delta}}\right)$ matches that of the exponential weights upper bound
- Optimal leading order term for $N = 2$
- Also give a nonasymptotic guarantee $\hat{u}(0) - E \leq v_{a^h}(0)$
- The discretization error E is computed explicitly and is $O\left(N\sqrt{N} \wedge \sqrt{N}\left(1 + \log\frac{1}{\delta}\right)\right)$