New Potential-Based Bounds for Prediction with Expert Advice

Vladimir A. Kobzar[1], joint work with 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 $q_t \in [-1,1]^N \sim a_t$, player's choice of expert $I_t \sim p_t$, these samples are revealed to both parties.

Preliminaries

- Starting time $T < 0$, final time $t = 0$
- *Instantaneous regret* (vector): $r = q_I \mathbb{1} - q$
- *Realized regret* at t (vector): $x = \Sigma_{\tau<t} r_\tau$
- *Final-time regret* (scalar) $R_T(p,a) = \mathbb{E}_{p,a} \max_i x_i$
- Player's objective to minimize, and adversary's objective is to maximize, R_T

Player value function

- Player p is Markovian: depends only on x, t
- *Value function*: v_p = expected final-time regret achieved by p if the game starts at realized regret x and time t and the adversary behaves optimally.

$$\begin{cases} v_p(x,0) = \max_i x_i & \text{(1a)} \\ v_p(x,t) = \max_a \mathbb{E}_{a,p}\, v_p(x+r,t+1), t<0 & \text{(1b)} \end{cases}$$

Intuition

- Value of a strategy is characterized by a dynamic program
- It is a discretization of a PDE, which captures the leading order behavior

References & acknowledgements

`http://proceedings.mlr.press/v125/kobzar20a/kobzar20a.pdf`; NSF grant DMS-1311833; Moore-Sloan Data Science Environment at NYU

Upper bound potential

A function w, nondecreasing in x_i, which solves

$$\begin{cases} w_t + \frac{1}{2} \max_{q\in[-1,1]^N} \langle D^2 w \cdot q, q\rangle \le 0 & \text{(2a)} \\ w(x,0) \ge \max_i x_i & \text{(2b)} \\ w(x + c\mathbb{1}, t) = w(x,t) + c & \text{(2c)} \end{cases}$$

- The associated player $p = \nabla w$
- Leads to an upper bound $v_p \le w$
- Bounds regret above: $v_p(0,T) = \max_a R_T(a,p)$
- *Exponential weights*: $w^e(x,t) = \Phi(x) - \frac{1}{2}\eta t$ where $\Phi(x) = \frac{1}{\eta}\log(\Sigma_i e^{\eta x_i})$ satisfies (2)

Proof of $v_p \le w$: step 1

Controlling the increase of w

- *Due to Δx*: (2c) implies $D^2 w\, \mathbb{1} = 0$ and by Taylor's thm,

$$\mathbb{E}_{p,a}\, w(x+r,t+1) - w(x,t+1)$$
$$\le \max_{q\in[-1,1]^N} \frac{1}{2}\langle D^2 w\cdot q, q\rangle \;[\le \;\eta/2 \text{ for } w^e]$$

where the choice of $p = \nabla w$ eliminated 1st-order term: $p \cdot q - \nabla w \cdot q = 0$
- *Due to Δt*:

$$w(x,t+1) - w(x,t) = w_t \;[= -\eta/2 \text{ for } w^e]$$

- By (2a), $\max_a \mathbb{E}_{p,a}[w(x+r,t+1)] - w(x,t) \le 0$

Proof of $v_p \le w$: step 2

Show $v_p \le w$ by induction

- Initialization: $v_p(x,0) \le w(x,0)$ by (1a) and (2b)
- Hypothesis: $v_p(x+r,t+1) \le w(x+r,t+1)$

$$\begin{aligned} w(x,t) &\ge \max_a \mathbb{E}_{p,a}\, w(x+r,t+1) && \textbf{[by step 1]} \\ &\ge \max_a \mathbb{E}_{p,a} v_p(x+r,t+1) && \textbf{[by hypothesis]} \\ &= v_p(x,t) && \textbf{[by (1b)]} \end{aligned}$$

Exp: $w^e(0,T) = \frac{1}{\eta}\log N + \frac{1}{2}\eta|T| = \sqrt{2|T|\log N}$ with $\eta = \sqrt{\frac{2\log N}{|T|}}$

Our contributions

- Potential-based viewpoint extends to adversaries, leading to lower bounds
- Upper and lower regret bounds $\equiv$ super and sub-solutions of certain PDEs
- Guidance for new strategies/improved bounds

Lower bound potential

- Adversary a is Markovian & "balanced" $\mathbb{E}_a q_i = \mathbb{E}_a q_j$
- Value function v_a for this adversary has a DP characterization similar to v_p
- *Lower bound potential* defn is similar to that of upper bound potential—a function u which solves

$$\begin{cases} u_t + \frac{1}{2}\mathbb{E}_a\langle D^2 u \cdot q, q\rangle \ge 0 & \text{(3a)} \\ u(x,0) \le \max_i x_i & \text{(3b)} \\ u(x + c\mathbb{1}, t) = u(x,t) + c & \text{(3c)} \end{cases}$$

- Since a is balanced, the 1st-order term is zero:

$$\mathbb{E}_{p,a}[q_I - \nabla u \cdot q] = \langle p - \nabla u, \mathbb{E}_a q\rangle = 0$$

- We used $\nabla u \cdot \mathbb{1} = 1$ by (3c) and $p \cdot \mathbb{1} = 1$
- $u \le v_a$ (modulo error E from higher order terms)
- Lower bound $u(0,T) - E(T) \le v_a(0,T) = \min_p R_T(a,p)$

Heat-based adversary a^h

- Gives best known leading-order prefactor $u(0,T) = \sqrt{-2\kappa T}\,\mathbb{E}_G \max G_i$ where $G \sim N(0,I)$,

$$\kappa = \begin{cases} 1 & \text{if } N = 2 \\ \frac{1}{2} + \frac{1}{2N} & \text{if } N \text{ is odd} \\ \frac{1}{2} + \frac{1}{2N-2} & \text{otherwise.} \end{cases}$$

- *Heat-based adversary* $a^h = \mathrm{Unif}(S)$ where

$$S = \begin{cases} \{q \in \{-1,1\}^N \mid \Sigma_{i=1}^N q_i = \pm 1\} & \text{for } N \text{ odd} \\ \{q \in \{-1,1\}^N \mid \Sigma_{i=1}^N q_i = 0\} & \text{for } N \text{ even} \end{cases}$$

- As lower bound potential, use the sol'n of the heat equation with κ as above

$$u_t + \kappa\Delta u = 0; \;\; u(x,0) = \max_i x$$

Heat-based adversary a^h vs. ML lit

- We provide a nonasymptotic guarantee $E(T) = O(N\sqrt{N} \wedge \sqrt{N \log N} + \sqrt{N}\log|T|)$
- a^h is asymptotically optimal for $N = 2$
- For large $|T|$, a^h gives a tighter l.b. than the previous state-of-the-art adversary a^s

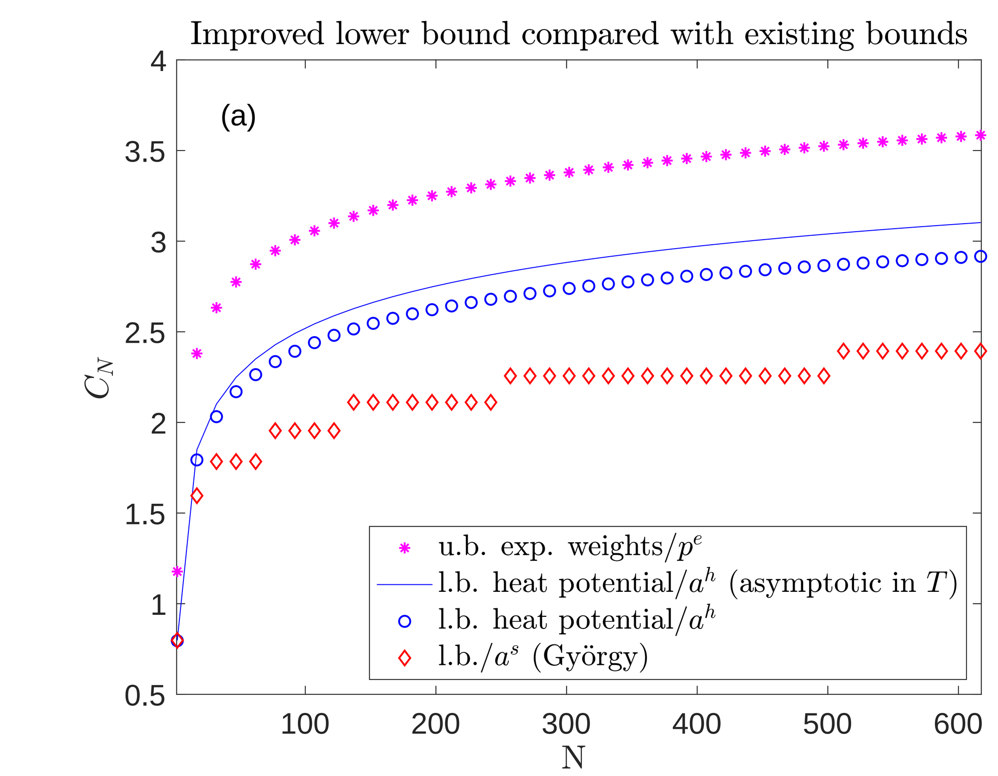

Figure 1:For an adversary a, $C_N\sqrt{|T|} \le \min_p R_T(a,p)$, and C_N determined for $|T| = 10^7$

New max potential

- The *max potential* is the explicit classical sol'n of

$$u_t + \kappa \max_i \partial_i^2 u = 0; u(x,0) = \max_i x$$

- Asymptotically optimal for $N = 2, 3$
- For small N and large $|T|$, max player p^m outperforms Exp

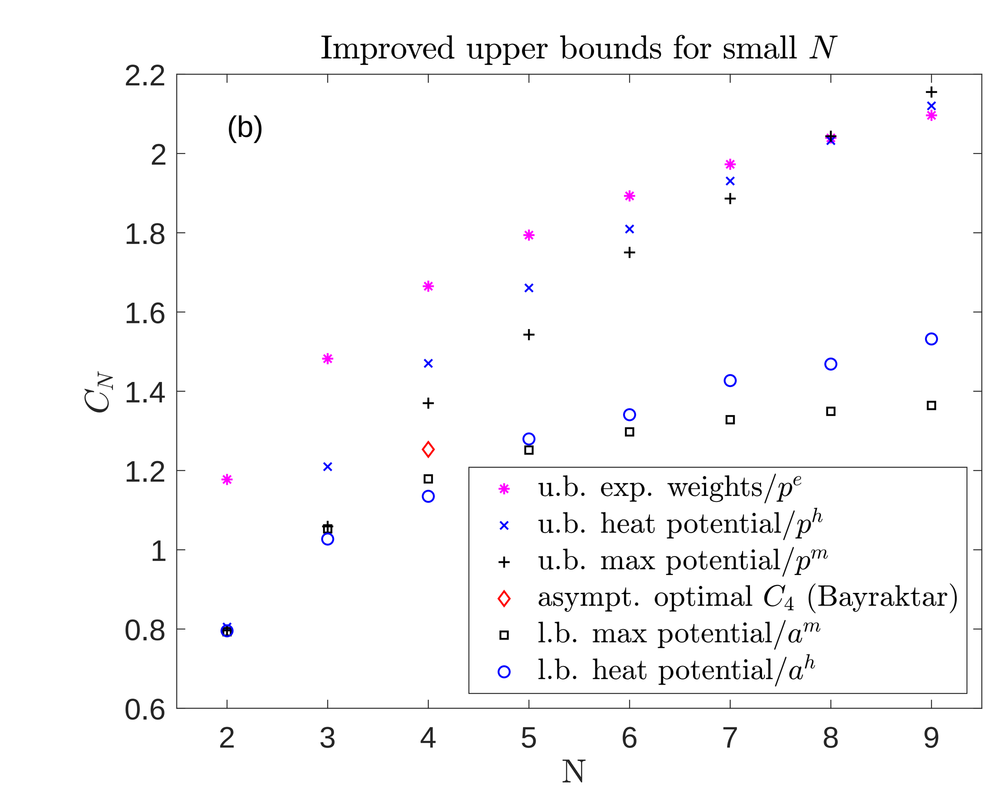

Figure 2:For a player p, $\max_a R_T(a,p) \le C_N\sqrt{|T|}$, and C_N determined for $|T| = 10^7$.