Complementi di Algoritmi e Strutture Dati

Bidding optimization in first-price auctions

Docente: *Nicolò Cesa-Bianchi* versione del 2 giugno 2026

We use Hedge/Exp3 to optimize bidding in first-price auctions. The buyer's revenue in a first-price auction is computed as follows: if the bid B is not larger than the highest competing bid M, then the item is sold to the highest bidder and the buyers's revenue is zero. Otherwise, if B is bigger than M, then the item is sold to the buyer with a profit of $V - B$ where V is the buyer's valuation for the item. Formally, the buyer's revenue is

$$g\big(B, M, V\big) = (V - B)\mathbb{I}\{B \geq M\}$$

We assume all quantities be in the unit interval $[0, 1]$ and we write $g_t(B) = (V_t - B)\mathbb{I}\{B \geq M_t\}$ for the buyer's revenue at time t as a function of the bid B given V_t and M_t, see Figure 1 for a pictorial representation.

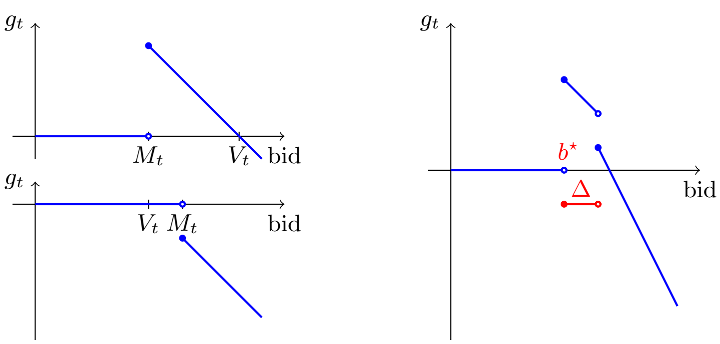

Figura 1: The revenue function is not Lipschitz. If $M_t \leq V_t$ (top left plot), then g_t is one-sided Lipschitz; conversely, if $M_t \geq V_t$ (bottom left plot), then g_t is still one-sided Lipschitz, but from the other side. Summing up the two types of revenues results in a function that may not be one-sided Lipschitz (right plot, where the two revenues of the other two plots are summed up. There, $b^\star$ is the optimal bid and Δ is the neighborhood of $b^\star$ where the total revenue is good enough).

We consider the following feedback: M_t is always observed after the auction is concluded, while V_t is only known if the auction is won, i.e., when $B_t \geq M_t$.

The right plot in Figure 1 shows that the optimal bidding region for $(V_1, M_1), \ldots, (V_T, M_T)$ may correspond to a very narrow interval Δ, outside of which the cumulative revenue is significantly smaller. This observation can be used to force any bidding algorithm to incur linear regret $\Omega(T)$. To overcome this problem, we assume the sequence of highest competing bids M_t is perturbed by a small amount of randomness.

A distribution μ on $[0, 1]$ is said to be σ-smooth for some $0 < \sigma \leq 1$ if for all intervals $[a, b] \subseteq [0, 1]$, we have $\mu([a, b]) \leq \frac{b-a}{\sigma}$. Note that for $\sigma = 1$ the only σ-smooth distribution is the uniform distribution.

Lemma 1 (Lipschitzness) *Let M_t be a σ-smooth random variable in $[0,1]^2$. Then the induced expected revenue function $\mathbb{E}[g_t]$ is $\frac{2}{\sigma}$-Lipschitz in $[0,1]$:*

$$\Big|\mathbb{E}\big[g_t(y) - g_t(x)\big]\Big| \le \frac{2}{\sigma}|y-x| \qquad \forall x,y \in [0,1]$$

DIMOSTRAZIONE. Let $x > y$ be any two bids in $[0,1]$, we have:

$$\begin{aligned}\Big|\mathbb{E}\big[g_t(x) - g_t(y)\big]\Big| &= \Big|\mathbb{E}\big[(V_t - x)\mathbb{I}\{M_t \le x\} - (V_t - y)\mathbb{I}\{M_t \le y\}\big]\Big| \\ &= \Big|\mathbb{E}\big[(V_t - x)\mathbb{I}\{y < M_t \le x\} - (x-y)\mathbb{I}\{M_t \le y\})\big]\Big| \\ &\le \mathbb{P}\big(M_t \in [y,x]\big) + (x-y) \le \tfrac{2}{\sigma}(x-y)\end{aligned}$$

□

Algoritmo 1 (Exp3-FPA)

Input: Exploration parameter $0 < \gamma \le 1$, finite grid $\mathcal{X} \subset [0,1]$ where $1 \in \mathcal{X}$.

1: Set uniform distribution p_1 over $\mathcal{X}$;
2: **for** $t = 1, 2, \ldots$ **do**
3: compute distribution $q_t(x) = (1-\gamma)p_t(x) + \gamma\mathbb{I}\{x = 1\}$ for $x \in \mathcal{X}$;
4: draw $B_t \sim q_t$ and post bid B_t;
5: observe M_t and $V_t\,\mathbb{I}\{M_t \le B_t\}$
6: for each $x \in \mathcal{X}$, compute the estimated revenue

$$\widehat{g}_t(x) = \frac{(V_t - x)\mathbb{I}\{M_t \le x\}}{\sum_{y \ge M_t} q_t(y)}\mathbb{I}\{M_t \le B_t\}$$

7: for each $x \in \mathcal{X}$, compute the new probability assignment

$$p_{t+1}(x) = \frac{\exp\left(\gamma \sum_{s=1}^t \widehat{g}_s(x)\right)}{\sum_{y \in \mathcal{X}} \exp\left(\gamma \sum_{s=1}^t \widehat{g}_s(y)\right)}$$

8: **end for**

Teorema 2 *Let $\mathcal{X} \subset [0,1]$ be a finite set containing 1. Then, the regret of Exp3-FPA run with $0 < \gamma \le 1$ against the best fixed bid in $\mathcal{X}$ is*

$$\max_{x \in \mathcal{X}} \sum_{t=1}^T g_t(x) - \mathbb{E}\left[\sum_{t=1}^T g_t(B_t)\right] \le \frac{\ln|\mathcal{X}|}{\gamma} + (e-1)\gamma T$$

DIMOSTRAZIONE. Because $q_t(x) = (1-\gamma)p_t(x) + \gamma\mathbb{I}\{x = 1\}$, we have $\sum_{y \ge M_t} q_t(y) \ge \gamma$. This implies $\gamma\widehat{g}_t(x) \le 1$ for each $x \in \mathcal{X}$, and so

$$\exp(\gamma\widehat{g}_t(x)) \le 1 + \gamma\widehat{g}_t(x) + (e-2)\gamma^2\widehat{g}_t(x)^2$$

where we used $e^x \le 1 + x + (e-2)x^2$ for all $0 \le x \le 1$. Let

$$w_t(x) = \exp\left(\gamma \sum_{s=1}^{t} \widehat{g}_s(x)\right) \qquad \text{and} \qquad W_{t+1} = \sum_{x \in \mathcal{X}} w_t(x)$$

Then,

$$\frac{W_{t+1}}{W_t} = \sum_{x \in \mathcal{X}} p_t(x) \exp\big(\gamma \widehat{g}_t(x)\big) \le 1 + \sum_{x \in \mathcal{X}} p_t(x) \Big(\gamma \widehat{g}_t(x) + (e-2)\gamma^2 \widehat{g}_t(x)^2\Big)$$

which implies

$$\begin{aligned}
\ln \frac{W_{t+1}}{W_t} &\le \sum_{x \in \mathcal{X}} p_t(x) \Big(\gamma \widehat{g}_t(x) + (e-2)\gamma^2 \widehat{g}_t(x)^2\Big) \\
&\le \frac{\gamma}{1-\gamma} \sum_{x \in \mathcal{X}} q_t(x) \Big(\widehat{g}_t(x) + (e-2)\gamma \widehat{g}_t(x)^2\Big)
\end{aligned}$$

Let $\mathbb{E}_t[\cdot] := \mathbb{E}\big[\cdot \mid B_1, \ldots, B_{t-1}, M_1, \ldots, M_T\big]$ and $\mathbb{P}_t(\cdot)$ defined accordingly. Since

$$\sum_{y \ge M_t} q_t(y) = \mathbb{P}_t(M_t \le B_t)$$

we have $\mathbb{E}_t\big[\widehat{g}_t(x)\big] = g_t(x)$ and

$$\mathbb{E}_t\left[\sum_{x \in \mathcal{X}} q_t(x) \widehat{g}_t(x)\right] = \mathbb{E}\big[g_t(B_t)\big]$$

Moreover,

$$\begin{aligned}
\mathbb{E}_t\left[\sum_{x \in \mathcal{X}} q_t(x) \widehat{g}_t(x)^2\right] &\le \mathbb{E}_t\left[\sum_{x \in \mathcal{X}} q_t(x) \frac{\mathbb{I}\{x \ge M_t\}\mathbb{I}\{M_t \le B_t\}}{\Big(\sum_{y \ge M_t} q_t(y)\Big)^2}\right] \\
&= \mathbb{E}_t\left[\sum_{x \in \mathcal{X}} q_t(x) \frac{\mathbb{I}\{x \ge M_t\}}{\sum_{y \ge M_t} q_t(y)}\right] = 1
\end{aligned}$$

Let $K = |\mathcal{X}|$. For each $x \in \mathcal{X}$,

$$\begin{aligned}
\gamma \sum_{t=1}^{T} g_t(x) - \ln K = \mathbb{E}\left[\gamma \sum_{t=1}^{T} \widehat{g}_t(x)\right] - \ln K &= \mathbb{E}\big[\ln w_{T+1}(x)\big] - \ln K \\
&\le \mathbb{E}\left[\ln \frac{w_{T+1}}{W_1}\right] = \sum_{t=1}^{T} \mathbb{E}\left[\ln \frac{W_{t+1}}{W_t}\right] \\
&\le \frac{\gamma}{1-\gamma}\left(\mathbb{E}\left[\sum_{t=1}^{T} g_t(B_t)\right] + (e-2)\gamma T\right)
\end{aligned}$$

Multiplying both sides by $\frac{1-\gamma}{\gamma}$ and rearranging gives

$$(1-\gamma)\sum_{t=1}^{T} g_t(x) \leq \mathbb{E}\left[\sum_{t=1}^{T} g_t(B_t)\right] + \frac{1-\gamma}{\gamma}\ln K + (e-2)\gamma T$$

which, using $\sum_{t=1}^{T} g_t(x) \leq T$ and $\frac{1-\gamma}{\gamma}\ln K \leq \frac{\ln K}{\gamma}$ gives

$$\sum_{t=1}^{T} g_t(x) - \mathbb{E}\left[\sum_{t=1}^{T} g_t(B_t)\right] \leq \frac{\ln K}{\gamma} + (e-1)\gamma T$$

concluding the proof. □

We now combine Lemma 1 with Theorem 2.

Teorema 3 *Consider the problem of repeated bidding in first-price auctions where $V_1, \ldots, V_T$ are arbitrary and $M_1, \ldots, M_T$ are drawn from an arbitrary sequence of σ-smooth distribution. Then, for all $T \geq (\ln K)/(e-1)$ there exists a learning algorithm such that*

$$\max_{b\in[0,1]} \sum_{t=1}^{T} g_t(b) - \mathbb{E}\left[\sum_{t=1}^{T} g_t(B_t)\right] \leq 6\left(\frac{1}{\sigma} + \sqrt{\ln T}\right)\sqrt{T}$$

Dimostrazione. Consider algorithm Exp3-FPA on the uniform grid $\mathcal{X}$ of $K = \lceil\sqrt{T}\rceil + 1$ bids. Then

$$\max_{b\in[0,1]} \mathbb{E}\left[\sum_{t=1}^{T} g_t(b)\right] \leq \max_{x\in\mathcal{X}} \mathbb{E}\left[\sum_{t=1}^{T} g_t(x)\right] + \frac{2}{\sigma}\gamma T \qquad \text{(using Lemma 1)}$$

$$\leq \mathbb{E}\left[\sum_{t=1}^{T} g_t(B_t)\right] + \frac{2}{\sigma}\gamma T + \frac{\ln K}{\gamma} + (e-1)\gamma T \qquad \text{(using Theorem 2)}$$

We obtain the claimed result by choosing $\gamma = \sqrt{\frac{\ln K}{(e-1)T}}$. □