Learning to bid in revenue-maximizing auctions

Thomas Nedelec, Noureddine El Karoui, Vianney Perchet

Thirty-sixth International Conference on Machine Learning

Some historical reminders

Roger Myerson

Revenue-maximizing auction

- if bidders are symmetric, second-price auction with well-defined reserve price is a revenue-maximizing auction.
- if we denote by F the CDF (f the PDF) of the value distribution of one bidder, the monopoly price r^* satisfies: $r^* = \frac{1-F(r^*)}{f(r^*)}$.
- For assymetric bidders, allocation based on the virtual value. Several approximations of the Myerson auction: eager/lazy, boosted second price, T-auctions, deep learning for auctions...

What is happening in practice : the online advertising use case

1. key assumption of Myerson : the auctioneer knows the value distribution F of the bidders : F is common knowledge.
2. in practice, this is not true...!
3. however, the auctioneer receives every day billions of bids of the different bidders : if the bidders bid truthfully, the auctioneer can learn F assuming bids are IID examples of the valuations of the bidders.

An example on Criteo Data

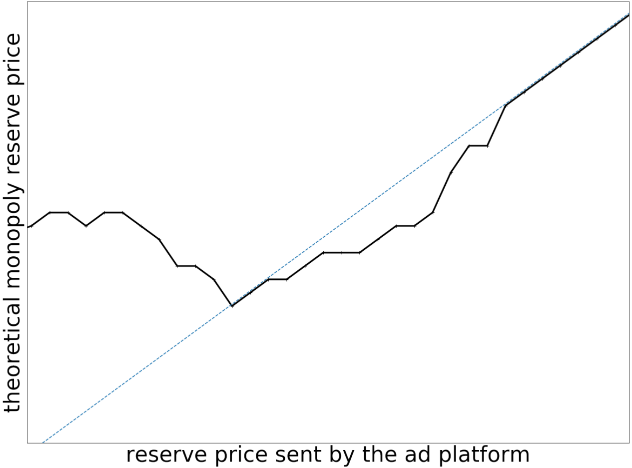

Figure: This plot was done on Criteo data. We bucketize all the requests we receive by the reserve price that was sent by a large ad platform. We then look on each bucket what would have been the optimal reserve price for Criteo. The plot is in log scale.

Key questions: the bidder's point of view

- Is it still dominant to bid truthfully when the seller is learning the reserve price from past bids ?
- What are the best bidding strategies when auctioneers are learning on past examples of bids to set the correct reserve price ?

A variational approach

Lemma

The utility of the strategic bidder using the strategy β increasing (ψ_B denotes the virtual value associated to the new distribution of bid) is given by:

$$\textit{Bidder Utility}(r) = \mathbb{E}_{X_i \sim F_i}\left((X_i - h_\beta(X_i)) G(\beta(X_i)) \mathbf{1}_{[X_i \geq x_\beta]} \right) .$$

with $h_\beta(X) = \psi_B(\beta(X)) = \beta(X) - \beta'(X)\frac{1-F(X)}{f(X)}$ *and* x_β *the reserve value.*

Experiments (exponential distribution)

Auction Type		K=2	K=3	K=4
Baselines	truthful revenue maximizing	0.30	0.24	0.21
	truthful welfare maximizing	0.50	0.33	0.25
Lazy second-price	Utility of strategic bidder	0.45 ± 0.001	0.31 ± 0.001	0.24 ± 0.001
	Uplift vs truthful bidding	+50%	+29%	+14%
Eager second-price	Utility of strategic bidder	0.52 ± 0.02	0.33 ± 0.02	0.25 ± 0.02
	Uplift vs truthful bidding	+73%	+37%	+19%
Myerson auction	Utility of strategic bidder	0.64 ± 0.001	0.45 ± 0.001	0.35 ± 0.001
	Uplift vs truthful bidding	+113%	+87%	+67%
Boosted second-price	Utility of strategic bidder	0.48 ± 0.03	0.41 ± 0.001	0.32 ± 0.001
	Uplift vs truthful bidding	+60%	+71%	+52%

Table: All bidders have an exponential value distribution with parameter $\lambda = 1$.