Sahil Dev
Dr. John Dickerson
University of Maryland, College Park
HONR269L

Game Theory, Auctions, and Deep Learning

Abstract

Game Theory is the analytical study of entities acting in a system to maximize their own gain. We go into detail about the core concepts of game theory. In particular, we discuss theoretical auction mechanism design, including a characterization of ideal auctions that we will apply to various auction situations. We also cover a deep learning model previously introduced by Paul Dütting et al. that is capable of producing auction results that approach the theorized optimal results by a number of measures, including auction revenue, while maintaining a number of additional desirable attributes. We also touch on potential modifications to the aforementioned deep learning architecture, which could theoretically decrease training time and increase robustness.

What is Game Theory?

Game Theory: the branch of mathematics concerning the analysis of strategic interaction between competitive participants

Nash Equilibrium: a stable point at which no single player in a game can change a single action and gain from it

Dominant Strategy: a strategy for a single player that is the best possible strategy no matter what any other player chooses to do.

Utility: the amount a player gains (opposite of **Loss**)

Payoff Matrix: shows the utility for each player under each strategy

Let's look at an example:

Suppose we have N countries subject to pollution. Each country has the choice to control their pollution production.

If a country C chooses not to control pollution, C does not have to pay expenses to do so, but every country, including C, must then spend $1 dealing with the consequences of pollution.

If C chooses to control pollution, C must pay $3, but each country then does not have to spend the $1 to deal with it.

1. If every country controls pollution, every country pays $3.

2. If no country controls pollution, every country pays $N. If we suppose N is much larger than 3, this is far from the ideal solution. Let's analyze it.

Consider the first situation. For any country C, C must pay $3 to control pollution. Suppose instead C refuses to control pollution production. C must now pay only $1 to handle the pollution it is producing. C's utility has increased by $2, at the expense of every other country increasing its loss by $1. It turns out every country can increase its utility (decrease its loss) by $2 by choosing to do the same thing. We have suddenly gone from the ideal first situation to the worst case second situation. Now every country must pay $N.

Suppose C considers paying the $3 to control pollution. Then C must pay $(N + 2) in total, increasing C's loss by $2. We have reached a Nash Equilibrium: no country benefits by changing their strategy. In fact, we have shown that every country's dominant strategy is to not control pollution.

Motivation

Auctions are a fundamental, inseparable part of game theory that come up everywhere in human behavior. As we will discuss later, even a system as simple selling an item can be characterized as an auction with one item and one bidder. Hence, any analysis of auction design has the capacity to have an incredible impact on economics. Designing optimal auction mechanisms is a notoriously difficult task requiring the construction of an allocation rule and a payment rule, and we still do not have a full construction for these rules in a 2-bidder, 2-item auction. We can, however, create increasingly good approximations for these rules using such techniques as deep learning. By developing models, we can improve performance and gain insight into ideal theoretical mechanisms.

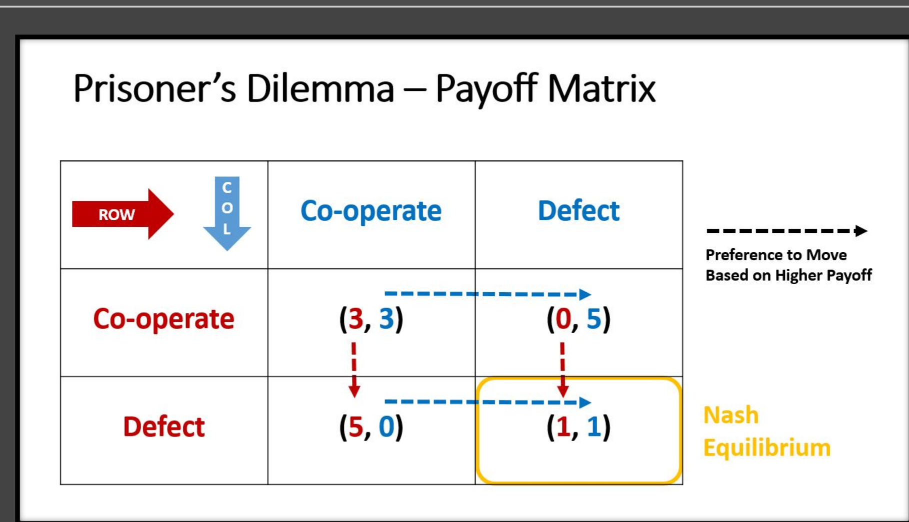

Definitions

Auction: a model in game theory consisting of m items, n bidders, and an auctioneer whose responsibility is to assign the allocation and pricing of items dependent on the bids

$$A = (X, p)$$

where X is an allocation rule and p is a payment rule

Bid: a function for the amount a bidder claims to value a set of items

$$b_i: \{0, 1\}^m \rightarrow \mathbb{R}^{\geq 0}, b_i \in V_i, i \in \mathbb{N}[1, n]$$

where V_i is the set of all possible bid/valuation functions for bidder i

Valuation: a function for the amount a bidder truly values a set of items

$$v_i: \{0, 1\}^m \rightarrow \mathbb{R}^{\geq 0}, v_i \in V_i$$

Allocation rule: a rule that assigns which items a bidder is allocated

$$X_i: V \rightarrow \{0, 1\}^m$$

where $V = \Pi V_i = (V_1, V_2, ..., V_n)$

Payment rule: a rule that assigns the payment a bidder must pay

$$p_i: V \rightarrow \mathbb{R}^{\geq 0}$$

Social Surplus: the extent to which the expectation that the bidder that values an item the most gets it is satisfied

$$S(X) = \Sigma v_i(X_i(b)), b \in V$$

Regret: the difference between the maximum possible utility of a bidder and their realized utility

What Does an Ideal Auction Look Like?

Dominant Strategy Incentive Compatible (DSIC): every bidder gets the highest possible gain from the auction by bidding their exact valuation

$$v_i(X_i(v_i, b_{-i})) - p_i(v_i, b_{-i}) \geq v_i(X_i(b_i, b_{-i})) - p_i(b_i, b_{-i}) \; \forall b_i \in V_i$$

where $b_{-i} \in V_{-i}$ is the vector of all bids with bidder i's entry removed

Maximum Social Surplus (MSS): the allocation is assigned in a way that maximizes the total possible valuation of those items

$$X = \text{argmax}_X (S(X)) = \text{argmax}_X (\Sigma v_i(X_i(b)))$$

Polynomial Time: the algorithm computing allocation and payments runs in polynomial time with respect to the number of bidders and items

$$O(n^p m^q)$$

The Second Price Auction: N-Bidders, One-Item

Maximum Social Surplus: The Second Price Auction defines the allocation rule so the highest bidder gets the item, satisfying MSS.

Dominant Strategy Incentive Compatible: The payment rule satisfies DSIC by setting the payment to the second highest bid. Let's see why.

If a bidder sets their bid above their valuation, there are two possibilities:

1. Nothing changes.
2. The overbid changes the bidder's position from below the highest to the highest, and the bidder must pay the second highest bid, which otherwise would have been the highest bid, higher than the bidder's valuation. The bidder gets the item but must now pay more than their valuation for the item, leading to a negative utility. This is undesirable.

If a bidder sets their bid below their valuation, there are two possibilities:

1. Nothing changes.
2. The underbid changes the bid from the highest to below the highest. The bidder pays nothing but no longer gets the item. Otherwise the bidder would get the item at a price lower than their valuation; their utility has changed from positive to zero. This is also undesirable.

Polynomial Time: Below is the algorithm. It runs in O(n) time.

1. Collect all bids (i, b_i). → O(n)
2. Select the two highest bids $(i, b_i)_1$ and $(i, b_i)_2$ such that $b_{i1} \geq b_{i2}$. → O(n)
3. Allocate the item to bidder i_1 at the price b_{i2}. → O(1)

We have shown that the Second Price Auction satisfies DSIC, MSS, and polytime, therefore satisfying all ideal properties of an auction.

One-Item, One-Bidder

Maximum Social Surplus: To satisfy this, we must allocate the item to the highest (only) bidder, regardless of their bid.

Dominant Strategy Incentive Compatible: If we set the payment to the bid value, the bidder will underbid. If we charge a constant price, the payment may be higher than the bid. All we can do is charge nothing.

Revenue Maximization: If we drop MSS and permit not allocating the item, we can set a price p so that the bidder is not allocated the item if b < p, and if b ≥ p, the bidder is allocated the item for the price p. Here we have discovered that the ideal way to sell an item to an individual is to set a constant price for it.

Optimal Auctions Through Deep Learning

Formulation as a Learning Problem

To define a learning problem, we must first select a loss function to minimize. The original paper (Paul Dütting et al.) adopts the negated revenue subject to DSIC (zero regret) as the loss function.

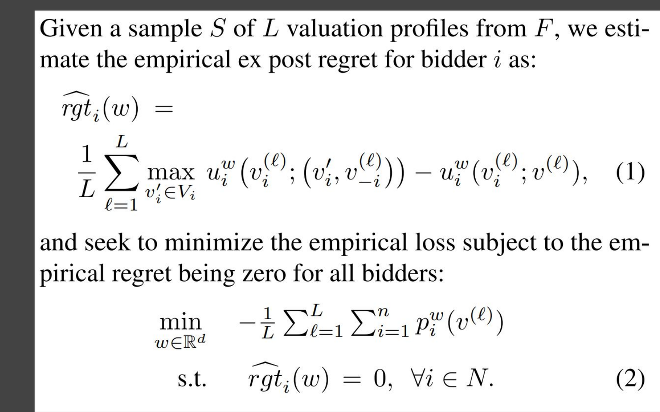
Given a sample S of L valuation profiles from F, we estimate the empirical ex post regret for bidder i as:

$$\widehat{rgt}_i(w) = \frac{1}{L}\sum_{\ell=1}^{L} \max_{v_i' \in V_i} u_i^w\left(v_i^{(\ell)}; \left(v_i', v_{-i}^{(\ell)}\right)\right) - u_i^w\left(v_i^{(\ell)}; v^{(\ell)}\right), \quad (1)$$

and seek to minimize the empirical loss subject to the empirical regret being zero for all bidders:

$$\min_{w \in \mathbb{R}^d} \; -\frac{1}{L}\sum_{\ell=1}^{L}\sum_{i=1}^{n} p_i^w(v^{(\ell)})$$
$$\text{s.t.} \quad \widehat{rgt}_i(w) = 0, \; \forall i \in N. \quad (2)$$

Neural Network Architecture and Performance

The architecture for RegretNet consists of an allocation network and a payment network, with 640,000 sample valuation profiles for training.

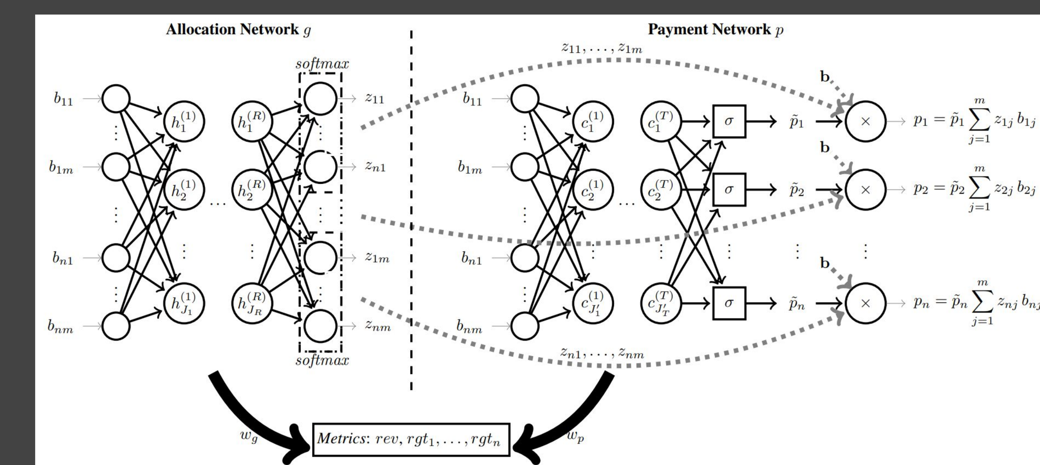

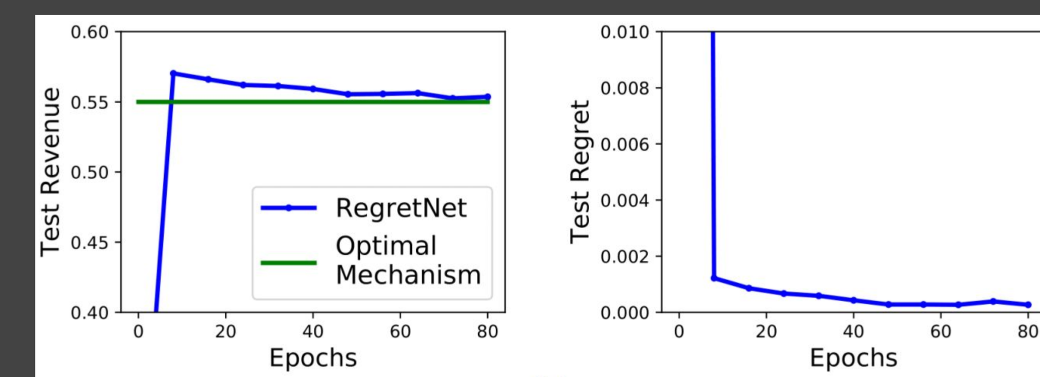

Possible Improvements

1. Change architecture from 2 hidden layers of 100 nodes (10^4 edges) to 3 layers of 10 nodes (300 edges) [9 hours → ≈ 20 minutes to train]
2. Use a Generative Adversarial Network (GAN) to train RegretNet (discriminator) and a generator that produces data to challenge the discriminator. This could improve convergence rate and robustness.

Sources

1. Nisan, Noam, et al. *Algorithmic Game Theory*. Cambridge University Press, 2007.
2. "Game Theory: Meaning of Game Theory by Lexico." *Lexico Dictionaries | English*, Lexico Dictionaries, www.lexico.com/definition/game_theory.
3. Dütting, Paul, et al. "Optimal Auctions Through Deep Learning." 2019.
4. Roughgarden, Tim. "Algorithmic Game Theory." 2013.
5. Bagchi, Shamit. "Prisoner's Dilemma - Payoff Matrix." *StudyCAS*, StudyCAS, 2015, studycas.com/images/PDGT.jpg.
6. "Generative Adversarial Networks | Google Developers." *Google*, Google, 8 Oct. 2019, developers.google.com/machine-learning/gan.