CS 6840 Algorithmic Game Theory Nov 4, 2024

Lecture 29: Limitations of Learning

Instructor: Eva Tardos *Scribe: Srisha Gaur*

1 Overview

We study first price multi-item auctions and show that it is NP-Hard to create an algorithm that gives small regret. More specifically, the difficulties in this problem lie within the algorithm itself rather than the learning aspect.

2 Background

Recall the First Price Multi-Item Auction setting discussed in previous lectures. We have m Players and n items. Define v_{ij} to be Player i's value on item j. In addition, the value to a Player for a set of items A is

$$v_i(A) = \max_{j \in A} v_{ij}.$$

Note that if a Player wins a group of items, they only gain value for a single item in that group.

Previously, we had shown that placing a bid of the form $\left(0, \ldots, 0, \frac{v_{ij}}{2}, 0, \ldots, 0\right)$ is learnable. That is, Player i bids $\frac{v_{ij}}{2}$ for item j and 0 for all other items. We showed that in this scenario the Price of Anarchy is $\frac{1}{2}$.

More generally, we showed that bidding of the form $(0, \ldots, 0, k\delta, 0, \ldots, 0)$, where δ allows us to discretize bids was learnable. The number of strategies, K, was the product of the number of items and $\frac{1}{\delta}$. This introduced a factor of K into the runtime of the algorithm and $\ln K$ in the regret, but there was not much impact.

3 Bidding Optimally

In general, most rational Players would not bid on a single item, as it does not seem to be a good idea.

This poses a new question. How much should a Player i bid? For example, they may not want to bid close to their values for all items, as they only receive utility from one item that is won. All other items that are bid on require the Player to pay a high value but receive no reward. In addition, for the item that is won, this Player receives little benefit.

A better idea may be to consider bidding ε rather than 0 all items except for one. On the remaining item j, the Player bids a shaded value. In the case that Player i does not bid enough on item j, there is still some small chance that they win a bid on a different item with little additional cost for the Player.

3.1 2-Player Bidding Strategy

Consider the first price multi-item auction in a two Player setting.

We let Player 1 have value $v_{1j} = V$ for all items j. In addition, they know that their opponent, Player 2, will either bid 1 or V', where $1 << V << V'$.

We also let Player 1 have full knowledge regarding Player 2's distribution. That is, Player 2 is given sets $T_1, \ldots, T_m$ and chooses a set T_i uniformly at random on which to bid value 1. On all items $j \notin T_i$ Player 2 bids V'.

Now assuming that Player 1 bids rationally, they will bid $1 + \varepsilon$, denoted 1^+, on some set S and 0 on all other items. Note that there is no reason to bid a value greater than 1^+, as it does not increase the probability that Player 1 bids and decreases the utility. Therefore, the question now becomes which is the optimal set for Player 1 to choose.

We first consider $u(S)$. That is, the expected utility obtained by Player 1 by bidding 1^+ on all items in set S. Note that the probability that a given set T_i is chosen is $\frac{1}{m}$. Player 1 would then win all items in $T_i \cap S$. Meaning, they receive value V if this intersection is non-empty. In addition, they must to pay 1^+ for these items. Therefore, our expected utility is

$$u(S) = \frac{1}{m} \sum_i \big(V \big| T_i \cap S \neq \emptyset\big) - |T_i \cap S|$$

3.2 Near Optimal Bids

Let $max_i |T_i| = d$ and $V = 2dm$. We claim that in this 2-player strategy an optimal or near optimal ensures that Player 1 will win an item, obtaining positive utility since V is large. In order to show this claim we provide and example strategy in which they win an item.

Suppose Player 1 bids 1^+ on the set $S = \cup_i T_i$. They obtain value V for the set of items won. The expected price paid is $\frac{1}{m} \sum_i |T_i|$. For each set T_i, the probability that Player 2 chooses this specific set to bid 1 on is $\frac{1}{m}$. The amount paid is then $1^+ \cdot |T_i| \approx |T_i|$. This gives us an expected utility of

$$\begin{aligned} u(S) &= V - \frac{1}{m} \sum_i |T_i| \\ &\geq V - d. \end{aligned}$$

Note that if a set S' does not guarantee a win,

$$u(S) \leq V \frac{m-1}{m} = V - 2d.$$

This follows from the fact that there the probability that $S' \cap T_i = \emptyset$ is at least $\frac{1}{m}$. In this case Player 2 receives a value of 0.

Observe that the utility obtained by bidding S is much greater than bidding S'. From this it follows that the optimal set S on which to bid must guarantee winning an item.

4 Hardness of Finding Optimal Set

4.1 Similar NP-Hard Problems

We now show that finding the optimal set S of items on which Player 1 should bid 1^+ is NP-Hard by reducing to REGULAR HITTING SET.

First recall the NP-Hard optimization problem HITTING SET. Given sets $X_1, \ldots X_m$, find the minimum set $H \subseteq \cup_i X_i$ such that $H \cap X_i \neq \emptyset$ for all i. Note that this is not exactly our problem. Rather than minimizing a set $|S|$, we must minimize $\sum_i |T_i \cap S|$.

We instead refer to an NP-Hard variant of this problem, REGULAR HITTING SET. Here, each set X_i has size d for all i. Similarly, each element is in exactly r sets, meaning $|\{i : j \in X_i\}| = r$. We depict this using a regular bipartite graph as shown below. The left side refers to the sets and the right side refers to the elements. In this case $|X_i| = 4$ and $|\{i : j \in X_i\}| = 2$.

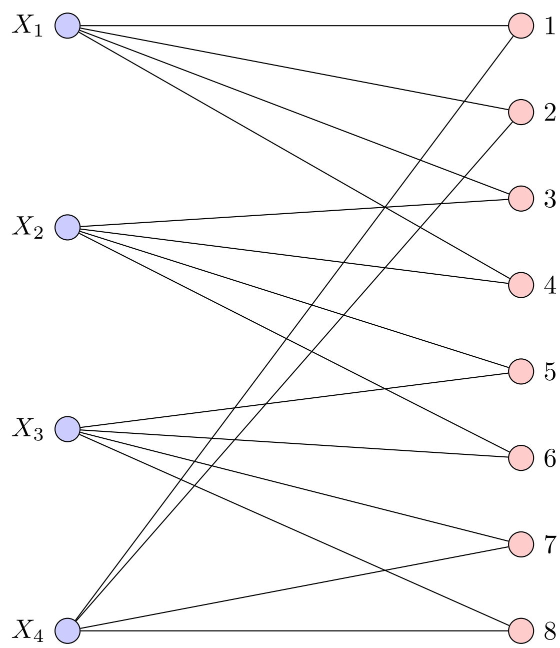

The above problem is NP-Hard to approximate this problem to a factor better than $\frac{1}{2} \ln r$.

4.2 Reducing to Bidding Problem

In order to show that finding the optimal bidding set S is NP-Hard, we reduce from REGULAR HITTING SET.

Our construction is as follows. Let X_i be the sets T_i that Player 2 can bid 1 on. Our hitting set H, then corresponds to the set S that Player 1 bids value 1^+. Therefore, each set T_i has size d and each element j is in exactly r sets T_i. In addition, we let m be an arbitrary constant and $V = 2dm$ to follow our example from Section 3.2.

We will show one direction of the reduction. Note that the other direction follows an analogous proof, so this suffices for showing both directions. That is, we show that if a set H is an optimal hitting set, then the corresponding set S maximizes utility for Player 1.

Since H is a hitting set and $S = H$, it must hold that $S \cap T_i \neq \emptyset$ for all i. Therefore, Player 1 will always win an item. This gives us an expected utility of

$$u(S) = V - \frac{1}{m}\sum_i |S \cap T_i|.$$

Note that the summation on the RHS is a summation over all sets and all items in these sets. That is,

$$V - \frac{1}{m}\sum_i |S \cap T_i| = V - \frac{1}{m}\sum_i \sum_{j \in T_i} 1.$$

By definition of a regular hitting set, each item j is included in exactly r terms in the above summation. We can therefore reverse the order of the summations to get

$$\begin{aligned} u(S) &= V - \frac{1}{m}\sum_{j \in S} r \\ &= V - \frac{r|S|}{m}. \end{aligned}$$

Since V, r, m are constants, since $|S| = |H|$ and H is a minimum hitting set it follows that $|S|$ must also be minimized. Therefore, $u(S)$ is maximized.

Furthermore, our approximation guarantees for Regular Hitting-Set imply that a set S' that approximates the optimal set will have

$$|S'| \geq \frac{1}{2}\ln r|S|.$$

4.3 In the Context of Learning.

Observe that our reduction did not consider learning, as the distribution was given. Furthermore, learning will not be helpful in solving this problem as it only discovers the distribution.

In addition, Follow the Leader as we just showed that choosing the best strategy is NP-Hard. Recall that regret can be written as $O(\sqrt{T \ln k})$, where K is the number of strategies.

Suppose their are n items and we discretize using a factor of δ. Then there are $\frac{1}{\delta}$ strategies for each item. This gives us $\left(\frac{1}{\delta}\right)^n$ strategies. Substituting for K in the expression for regret gives us

$$O(\sqrt{T \ln K}) = O\left(\sqrt{nT \ln \frac{1}{\delta}}\right).$$

Although the above regret is small, the issue lies in implementing a single step of Follow the Leader.

Recall that in Multiplicative Weights there were too many probabilities that must be maintained in order to run. Perturbed Follow the Leader has both the hardness properties of Follow the Leader and Multiplicative Weights.