CSCI699: Topics in Learning and Game Theory

Lecture October 09

Lecturer: Ilias Diakonikolas *Scribes: Kai Wang, Junyi Liu*

1 Mechanism Design with Samples or Learning Near-Optimal Auction

Recall the things we will need:

1. PAC Learning: There exists function $m_H : (\epsilon, \delta) \rightarrow N$ such that for all $\epsilon, \delta > 0$, distribution D over X, target concept $f \rightarrow \{0, 1\}$, then the ERM algorithm after $m_H(\epsilon, \delta)$ samples guarantees to output a hypothesis h with high probability $1 - \delta$ to be ϵ-approximate to the target concept. Here, we are going to use the similar idea. Without knowing the buyer's payoff functions and distribution, but according to the samples from the history data, the seller could learn the buyer's payoff function (or so called the target concept) in order to maximize the revenue.

2. Myerson's Optional Auction

2 Revenue Maximization

Given a bunch of users with multiple products, the seller wants to maximize the revenue. But the Seller doesn't know the payoff function. He only knows the samples draw from the distribution D of the payoff function. Given the values of bidder, maximize the expectation of revenue.

There may be some closed forms of calculating

Example 1 (Single Item, single bidder)**.** *There is only one buyer, one item, one seller. Suppose the seller is using the posted price auction.*

Posted price auction of the reserve price $v \in [0, 1], v \sim D$

$$r(p, v) = \begin{cases} p, v \geq p \\ 0, otherwise \end{cases}$$

The expected revenue would be

$$R_D(p) = \mathbb{E}_{v\sim D}\ r(p,v) = \mathbb{E}_{v\sim D}\ [p \cdot \mathbf{1}_{v\geq p}] = p \cdot P(v \geq p) = p(1-F(p))$$

where function F is the cumulative function of distribution D, i.e., $F(p) = P(v \leq p|v \sim D)$.

In practice we don't know the distribution D explicitly. But we could observe the price auctions $\{v_t\}$ along the time. Suppose that these samples are drawn from the distribution D but have no prior information about the distribution itself. It is possible that all the samples are drawn from the same point, i.e. m completely identical sample. In this case, we would never be able to learn the distribution. Therefore we would never be 100% sure that our algorithm provides a correct approximated concept. So, our task is to guarantee to learn the payoff function (distribution) with high probability.

Assumption 2. *We have access to IID samples from D.*

Given a sample S of size m, we want to compute a reserve price p_S, s.t. with high probability (ex. $\frac{9}{10}$), we have

$$R_D(p_S) \geq \max_{p\in[0,1]} R_D(p) - \epsilon. \tag{1}$$

in other words, we want the reserve price p_S to achieve ϵ-approximated optimal value with high probability.

The approximation is highly related to the number of samples we observed. Thus the question becomes what the minimum m is so that the above inequality holds with high probability? Such m can be interpreted as the sample complexity of revenue maximization. This idea is similar to the idea in PAC learning, which is the number of the samples we need to guarantee the approximation with high probability and small error rate. The answer is $m = \Theta(\frac{1}{\epsilon^2})$.

In Example 1, suppose we can approximate the cumulative function $F(p)$ by $\hat{F}(p)$ with $|\hat{F}(p) - F(p)| < \frac{\epsilon}{p}$ in high probability. Let $p = \arg\max_{p'\in[0,1]} p'(1-F(p'))$ be the true optimal solution for maximizing revenue and $p_S = \arg\max_{p'\in[0,1]} p'(1-\hat{F}(p'))$ be the reserve price calculated by using the approximated cumulative function $\hat{F}(p)$.

$$R(p) - R_D(p_S) = (p(1 - F(p))) - (p_S(1 - \hat{F}(p_S))) \tag{2}$$
$$= (p(1 - \hat{F}(p)) - p_S(1 - F(\hat{p_S}))) + p(\hat{F}(p) - F(p)) \tag{3}$$
$$\leq 0 + p\frac{\epsilon}{p} \tag{4}$$
$$= \epsilon \tag{5}$$

The first term in the Inequality 3, 4 is because $p_S = \arg\max_{p' \in [0,1]} p'(1 - \hat{F}(p'))$ optimizes the approximated cumulative function term, thus the value will be greater than or equal to the value of using any other $p \in [0, 1]$.

Therefore, it suffices to show that we can approximate the unknown cumulative function $F(p)$ with high probability by using m samples draw from the distribution. Determining how many samples m do we need becomes the main problem in this example.

3 Rademacher Complexity and Covers

The first method we will use to derive the sample complexity is Radmemacher Complexity + Covers. We first give some definitions as follows.

Definition 3. *Consider a set of samples $S = \{v_1, \ldots, v_m\}$, a hypothesis class $\mathcal{H}$, and the loss function $l(h, v)$ for $h \in \mathcal{H}$ and $v \in S$. Let*

$$R(S, \mathcal{H}) := \mathbb{E}_\sigma \left[\sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{t=1}^{m} \sigma_t \cdot l(h, v_t) \right] \tag{6}$$

where $\sigma = (\sigma_1, \cdots, \sigma_m) \in \{\pm 1\}^m$ and $P(\sigma_i = 1) = P(\sigma_i = -1) = \frac{1}{2}$ for any $i = 1, \ldots, m$.

Let S be the sample set of the previous auction values, $[0, 1]$ be the hypothesis class of the reserve price $\mathcal{H}$, $r(p, v_t)$ be the loss function for $p \in [0, 1]$ and $v_t \in S$. Then we have $R(S, H) = \mathbb{E}_\sigma \left[\sup_{p \in [0,1]} \frac{1}{m} \sum_{t=1}^{m} \sigma_t \cdot r(p, v_t) \right]$

3.1 Recall the Lecture 2

Recall the definition ϵ-representative

Definition 4 (Definition 9 of the Lecture 2)**.** *A sample* $S = \{z_1, z_2, ..., z_m\}$ *is* ϵ-*representative (w.r.t domain* Z*, hypothesis class* H *and loss function* $l(h,z)$*) if*

$$\sup_{h \in \mathcal{H}} |L_D(h) - L_S(h)| \leq \epsilon$$

where $L_D(h) = \mathbb{E}_{z \sim D}[l(h,z)]$ *and* $L_S(h) = \mathbb{E}_{z \sim \mathcal{U}(S)}[l(h,z)]$ *(*$\mathcal{U}(S)$ *is the uniform distribution over* S*.)*

Definition 5. $Rep_D(F,S) = \sup_{f \in F} |L_D(f) - L_S(f)|$*, where* $F = \{f_h : h \in \mathcal{H}\}, f_h(z) = l(z,h)$*.*

$Rep_D(F,S)$ provides a measure of the largest difference between expected true error and expected empirical error.

Definition 6. *Rademacher Complexity of* F *(w.r.t sample* S*) is defined as*

$$R_S(F) = \frac{1}{m}\mathbb{E}_{\sigma_i}[\sup_{f \in F} \sum_{i=1}^{m} \sigma_i f(z_i)]$$

where each σ_i *is an independent* $\{-1,1\}$ *coin flip.*

Lemma 7 (Lemma 10 of the Lecture 2)**.**

$$\mathbb{E}_{S \sim D^m}[Rep_D(F,S)] \leq 2\mathbb{E}_{S \sim D^m}[R_S(F)]$$

3.2 Estimation

Return to our problem, if p_S is the empirically optimal price, i.e.

$$p_S = \arg\max_{p \in [0,1]} \left[R_S(p) \equiv \frac{1}{m} \sum_{t=1}^{m} r(p, v_t) \right].$$

From the Lemma 7, we have

$$\mathbb{E}_S[Rep_D(H,S)] \leq 2\mathbb{E}_S[R(S,H)] \tag{7}$$

The LHS in the Lemma 7 is the expected largest difference between true error and empirical error. Here, we replace all the loss function by the revenue function. Thus the LHS can be expressed as the expected largest difference between true revenue $\sup_{p\in[0,1]} R_D(p)$ and the empirical revenue $R_D(p_S)$, which is

$$\mathbb{E}_S[\sup_{p\in[0,1]} |R_D(p) - R_D(p_S)|] = \sup_{p\in[0,1]} R_D(p) - \mathbb{E}_S[R_D(p_S)] \leq 2\mathbb{E}_S[R(S,H)]$$

which implies

$$\mathbb{E}_S[R_D(p_S)] \geq \sup_{p\in[0,1]} R_D(p) - 2\mathbb{E}_S[R(S,\mathcal{H})] \tag{8}$$

If we have $\mathbb{E}_S[R(S,\mathcal{H})] < \epsilon$, then the revenue achieved by the reserve price p_S based on the samples S is 2ϵ-approximated to the highest revenue.

Thus our goal is to bound from above of $\mathbb{E}_S[R(S,\mathcal{H})]$ such that $\mathbb{E}_S[R(S,\mathcal{H})] < \epsilon$. The idea of deriving such upper bound is to discretize the hypothesis class $\mathcal{H} = [0,1]$. Consider the revenue function $r(p,v) = \begin{cases} p, v \geq p \\ 0, \text{otherwise} \end{cases}$ in the example, let $\mathcal{H}_\epsilon = \{0, \epsilon, 2\epsilon, 3\epsilon, \ldots, 1\} \cup S$.

Claim 8. *$\forall p \in [0,1], \exists p_\epsilon \in H_\epsilon$ s.t. $\forall v_t \in S$, we have*

$$|r(p,v_t) - r(p_\epsilon, v_t)| \leq \epsilon$$

Proof. We first reorder the sample set such that $0 = v_0 \leq v_1 \leq \ldots \leq v_m \leq v_{m+1} = 1$. Suppose $p \in [v_{t-1}, v_t]$ for some $t \in [m+1]$. If the closest multiple of ϵ below p is in this interval, set p_ϵ to be this value, otherwise set $p_\epsilon = v_{t-1}$. Then we have $|r(p,v_t) - r(p_\epsilon, v_t)| \leq \epsilon$ for any $v_t \in S$. □

Remark 1. *In general case, given any revenue function r, if the revenue function is continuous and lies in a bounded interval, then we can use the inverse function r^{-1} to select the discretization points.*

Lemma 9. *Let $\mathcal{H}$ be any hypotheses class and S be the sample set, if $\mathcal{H}_\epsilon$ is ϵ-cover of S, then*

$$R(S,\mathcal{H}) \leq R(S,\mathcal{H}_\epsilon) + \epsilon$$

3.3 Hoeffding's Inequality

Combine Lemma 9 with inequality (8), our goal becomes to bound $\mathbb{E}_S[R(S, \mathcal{H}_\epsilon)]$ with $|H_\epsilon| \leq O(m + \frac{1}{\epsilon})$. Recall the Hoeffding's inequality

Lemma 10 (Hoeffding's inequality)**.** *Suppose* $\bar{X} = \frac{1}{m}(X_1 + X_2 + ... + X_m), X_i \in [a_i, b_i]$, *we have*

$$P(\bar{X} - \mathbb{E}[\bar{X}] \geq t) \leq exp(-\frac{2m^2t^2}{\sum_{i=1}^{m}(b_i - a_i)^2}) \tag{9}$$

Theorem 11. *Given* m *samples drawn from distribution* D*,* $R(S, H_\epsilon) < \epsilon$ *with probability* $1 - \delta$*, where* $m \geq \frac{\log \frac{1}{\epsilon}}{\epsilon^2}, \delta = |H_\epsilon| exp(-m\epsilon^2/2)$.

Proof. Given $p \in [0, 1]$, let $X_i = \sigma_i r(p, v_i)$, clearly we have $\mathbb{E}_{\sigma_i}[X_i] = \mathbb{E}_{\sigma_i}[\sigma_i r(p, v_i)] = 0$ and $X_i \in [-1, 1]$. By Hoeffding's inequality,

$$P(\bar{X} \geq \epsilon) = P(\bar{X} - \mathbb{E}[\bar{X}] \geq \epsilon) \leq exp(-\frac{2m^2\epsilon^2}{\sum_{i=1}^{m}(b_i - a_i)^2}) = exp(-m\epsilon^2/2) \tag{10}$$

Apply the union bound to all the reserve price in H_ϵ, we have

$$P(R(S, H_\epsilon) \geq \epsilon) = P(\sup_{p \in H_\epsilon} \bar{X}^p \geq \epsilon) \leq |H_\epsilon| P(\bar{X} \geq \epsilon) \leq |H_\epsilon| exp(-m\epsilon^2/2) \tag{11}$$

Let $\delta = |H_\epsilon| exp(-m\epsilon^2/2)$, i.e. $\epsilon = \sqrt{\frac{\log \frac{|H_\epsilon|}{\delta}}{m}}$, $m = \frac{\log \frac{|H_\epsilon|}{\delta}}{\epsilon^2}$

Suppose the above conditions of $m, \delta, \epsilon, |H_\epsilon|$ are satisfied, then with probability greater than $1 - \delta$, we have

$$R(S, \mathcal{H}_\epsilon) \leq \epsilon \tag{12}$$

which is exactly what we want.

The remaining is to check whether $m \geq \frac{\log \frac{|H_\epsilon|}{\delta}}{\epsilon^2} = \frac{\log \frac{m + \frac{1}{\epsilon}}{\delta}}{\epsilon^2}$ could be satisfied. And this could be achieved when $m \geq \frac{\log \frac{1}{\epsilon}}{\epsilon^2}$ □

3.4 Conclusion of Example 1

Now back to our example, we know $R_D(p) = p(1 - F(p))$. We have the following claim as we previously proved:

Theorem 12. *Suppose that we have $\hat{F}$ s.t.*

$$\sup_{p\in[0,1]} |F(p) - \hat{F}(p)| \leq \epsilon \tag{13}$$

Then $R_{\hat{D}}(p) \geq R_D(p) - \epsilon$

So it suffices to find such an $\hat{F}$. We construct an "Empirical CDF", i.e. $\hat{F}(z) = \frac{1}{m}\sum_{i=1}^{m} 1_{x_i \leq z}$. After some similar calculation, we have

$$\sup |\hat{F}(z) - F(z)| \leq \sqrt{\frac{10}{m}} \tag{14}$$

with probability $\frac{9}{10}$ (or $\sqrt{\frac{1}{m\epsilon}}$ with probability $1-\epsilon$). And this construction guarantees that after learning $m \geq O(\frac{1}{\epsilon^2})$ (or $m \geq \frac{\log\frac{1}{\epsilon}}{\epsilon^2}$) samples drawn from distribution D, with high probability $1-\delta$ we can achieve an ϵ-approximate to the maximum revenue.