CAS CS 591 B: Communication Complexity

Prof. Mark Bun Fall 2019

Lecture Notes 4:

Distributional Complexity, Yao's Principle

Reading.

- Rao-Yehudayoff Chapter 3, "Minimax"

We would like to develop techniques for proving lower bounds against randomized communication protocols. We saw one technique for doing this in the last lecture, which was to simulate a private-coin protocol with a deterministic protocol using an exponential blow-up in communication. In this lecture we will start to develop techniques which allow us to prove stronger lower bounds.

Recall that Newman's Theorem showed that public coins do not give us substantially more power than private coins. It is often easier to prove both upper and lower bounds against the public coin model. The technique we will use to prove public coin lower bounds relies on a different way of introducing randomness into the communication model: Instead of considering distributions over protocols, we consider distributions over inputs.

Definition 1. Let $f : X \times Y \to \{0, 1\}$ be a two-party function and let $\mu : X \times Y \to [0, 1]$ be a probability measure. The *ε-error distributional complexity of f with respect to μ*, denoted $D^{\mu}_{\varepsilon}(f)$, is the least cost of a deterministic protocol Π such that

$$\Pr_{(x,y)\sim\mu}[\Pi(x, y) = f(x, y)] \geq 1 - \varepsilon.$$

Example 2. Let μ be the uniform distribution over $\{0, 1\}^n \times \{0, 1\}^n$. Then $D^{\mu}_{1/4}(\mathrm{GT}_n) \leq 2$. Alice can send the most significant bit x_1 of her input, and Bob can check whether $x_1 > y_1$. This will give the correct answer on at least 3/4 of input pairs. (I.e., if $x_1 \neq y_1$, which happens with probability 1/2, this protocol always succeeds. Otherwise, if $x_1 = y_1$, this protocol succeeds with probability 1/2.)

To prove a public coin randomized communication lower bound, it suffices to identify a distribution μ for which D^{μ}_{ε} is large.

Theorem 3. *For every distribution $\mu : X \times Y \to [0, 1]$, we have*

$$\mathbf{BPP}^{pub}_{\varepsilon}(f) \geq D^{\mu}_{\varepsilon}(f).$$

Proof. Fix a distribution μ. We show that a public coin randomized protocol for computing f can always be converted into a deterministic protocol which succeeds with respect to μ. Let $\Pi(x, y; r)$ denote such a protocol. Since Π succeeds on every input (x, y), it in particular succeeds with respect to (x, y) drawn from μ:

$$\begin{aligned}\forall(x,y)\Pr_r[\Pi(x,y;r)=f(x,y)]&\geq 1-\varepsilon\\ \Rightarrow \Pr_{r,(x,y)\sim\mu}[\Pi(x,y;r)=f(x,y)]&\geq 1-\varepsilon\\ \Rightarrow \exists r \Pr_{(x,y)\sim\mu}[\Pi(x,y;r)=f(x,y)]&\geq 1-\varepsilon.\end{aligned}$$

Consider the deterministic protocol $\Pi(\cdot,\cdot;r)$ obtained by fixing the coins of Π to the random string guaranteed by the final inequality. Then this is a good protocol with respect to μ. □

The proof of Theorem 3 appears lossy. It seems we've replaced a universal quantifier over (x, y) with an existential quantifier over r. Nevertheless, it turns out we haven't lost anything at all, and in fact distributional complexity is equivalent to $\mathbf{BPP}^{pub}$ complexity:

Theorem 4.

$$\mathbf{BPP}_{\varepsilon}^{pub}(f)=\max_{\mu} D_{\varepsilon}^{\mu}(f).$$

One way to interpret this theorem is that *every* lower bound against public coin protocols can be obtained by identifying a suitable hard distribution for deterministic protocols. This equivalence is known as Yao's minimax principle. To describe the proof, we take a detour into game theory.

1 Zero Sum Games

In a two player game, a player Row holds a set of *strategies* R and a player Column holds a set of strategies C. To play the game, Row selects a strategy $r \in R$, Column (simultaneously) selects a strategy $c \in C$, and the players receive rewards $U_R(r, c)$ and $U_C(r, c)$ respectively. Such a game is *zero-sum* if the reward to Row is exactly the negative of the reward to Column, i.e., there exists a *payoff function* $U(r, c)$ such that $U_R = U$ and $U_C = -U$.

The goal of the Row player is to maximize $U(r, c)$ and the goal of the column player is to minimize $U(r, c)$. In addition to playing pure strategies, the players can play mixed strategies, which are distributions over pure strategies. Let $p \in [0, 1]^{|R|}$ be a vector with $\sum_{r\in R} p_r = 1$, denoted $p \in \Delta(R)$. Row playing the mixed strategy p corresponds to playing $r \in R$ with probability p_r. Similarly, Column playing a mixed strategy $q \in \Delta(C)$ corresponds to playing $c \in C$ with probability q_c. If Row and Column both play mixed strategies, the *expected* payoff to Row is

$$U(p,q) := \sum_{r\in R}\sum_{c\in C} p_r U(r,c) q_c.$$

(This is an abuse of notation. We could remove the abuse by replacing $U(r, c)$ with $U(e_r, e_c)$ where e_r and e_c are the indicator vectors for actions r and c, respectively.) Let's consider two variations of this game. In the first variation, we (apparently) tip the scales in favor of Column.

Variant 1. Row has to announce a mixed strategy p and then Column is allowed to best respond with a strategy q attaining payoff

$$\min_{q\in\Delta(C)} U(p,q).$$

In this variation of the game, Row should choose a mixed strategy p so as to *maximize* this quantity, so as to guarantee payoff

$$\max_{p\in\Delta(R)} \min_{q\in\Delta(C)} U(p,q).$$

Variant 2. Similarly, we can rig the game in favor of Row, and force Column to announce a mixed strategy q beforehand. In this case, Column should choose q so as to *minimize* $\max_{p\in\Delta R} U(p,q)$, guaranteeing payoff

$$\min_{q\in\Delta(C)} \max_{p\in\Delta(R)} U(p,q).$$

Forcing Row to go first should always make that player worse off, so we have

$$\max_{p\in\Delta(R)} \min_{q\in\Delta(C)} U(p,q) \leq \min_{q\in\Delta(C)} \max_{p\in\Delta(R)} U(p,q).$$

Von Neumann's Minimax Theorem asserts that this inequality is actually an equality. Each player can commit to an "equilibrium strategy" (p^*, q^*) such that, even if they have to announce this strategy first, they can guarantee a particular payoff called the "value" of the game.

Theorem 5 (Minimax Theorem)**.** *For every two-player zero sum game with payoff function U, there exists a* value

$$\text{val} = \max_{p\in\Delta(R)} \min_{q\in\Delta(C)} U(p,q) = \min_{q\in\Delta(C)} \max_{p\in\Delta(R)} U(p,q).$$

Example 6. Consider the game Rock-Paper-Scissors. We can conveniently organize the payoffs of pure strategies in the following matrix:

$$U = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}.$$

The equilibrium strategies for this game are for each player to play the uniform distribution over the rock, paper, and scissors pure strategies. If Row announces this strategy first, uniform is still an optimal strategy for Column and guarantees expected payoff zero. (Similarly if Column announces first.)

2 Proof of Theorem 4

Let $c = \max_\mu D^\mu_\varepsilon(f)$. We define a two-player zero sum game between a protocol designer (Row player) and an adversary (Column player). Their available strategies are as follows:

Protocol Designer (Row): Let P_c denote the set of all deterministic protocols with cost c. Pure strategies are $\Pi \in P_c$. Mixed strategies are all randomized protocols $\tilde{\Pi}$, which are distributions in $\Delta(P_c)$.

Adversary (Column): Pure strategies are inputs (x,y). Mixed strategies are distributions μ over $X \times Y$.

The payoff $U(\Pi, (x, y))$ to the protocol designer is 1 if $\Pi(x, y) = f(x, y)$ and 0 otherwise. That is, the Row player wins iff the protocol chosen succeeds on the input chosen by the adversary.

By assumption, for every mixed strategy μ for the Column player, there exists a (pure) strategy for the Row player with expected payoff $\geq 1 - \varepsilon$. This puts us in the setting of Variant 2: The Column player announces first, and the Row player gets to pick a strategy so as to maximize their payoff. That is,

$$\min_{\mu \in \Delta(X \times Y)} \max_{\tilde{\Pi} \in \Delta(P_c)} U(\tilde{\Pi}, \mu) \geq 1 - \varepsilon.$$

By the Minimax Theorem, this implies

$$\max_{\tilde{\Pi} \in \Delta(P_c)} \min_{\mu \in \Delta(X \times Y)} U(\tilde{\Pi}, \mu) \geq 1 - \varepsilon.$$

Interpreting this as the setting of game Variant 1, there exists a mixed strategy for the Row player such that for every (mixed, and in particular, pure) strategy of the Column player, Row can guarantee payoff $\geq 1 - \varepsilon$. This mixed strategy exactly corresponds to a protocol with success probability at least $1 - \varepsilon$ on every input (x, y).

3 "Proof" of Theorem 5

One way to prove the Minimax Theorem is to appeal to strong duality of linear programming. This is only a "proof" in the sense that the latter is already a deep result which we'll just take on faith, at least for the time being. But strong duality is itself important for understanding a number of concrete lower bound techniques, so we may as well start using it now.

Recall that our expected payoffs in the two variants of the game are

$$\begin{aligned} &\max_{p \in \Delta(R)} \min_{q \in \Delta(C)} U(p, q) && \text{if Row goes first} \\ &\min_{q \in \Delta(C)} \max_{p \in \Delta(R)} U(p, q) && \text{if Column goes first.} \end{aligned}$$

However, if Row goes first, observe that Column can always choose a *pure* strategy so as to minimize $U(p, c)$. (And similarly if Column goes first.) So we can actually write these payoffs as

$$\begin{aligned} &\max_{p \in \Delta(R)} \min_{c \in C} \sum_{r \in R} U(r, c) p_r && \text{if Row goes first} \\ &\min_{q \in \Delta(C)} \max_{r \in R} \sum_{c \in C} U(r, c) q_c && \text{if Column goes first.} \end{aligned}$$

We can encode the first quantity as the solution to the following linear program:

$$\begin{aligned} \max \quad & t \\ \text{s.t.} \quad & t - \sum_{r \in R} U(r, c) p_r \leq 0 && \forall c \in C \\ & \sum_{r \in R} p_r = 1 \\ & p_r \geq 0 && \forall r \in R \end{aligned}$$

The dual of this LP is given by:

$$\begin{aligned}
\min \quad & s \\
\text{s.t.} \quad & s - \sum_{c \in C} U(r,c) q_c \geq 0 & \forall r \in R \\
& \sum_{c \in C} q_c = 1 \\
& q_c \geq 0 & \forall c \in C
\end{aligned}$$

By strong LP duality, the value of the dual LP is equal to the value of the primal LP. But the dual LP can be seen to exactly capture the second quantity, proving the Minimax Theorem.