Privacy in Statistics and Machine Learning — Spring 2025
In-class Exercises for Lecture 20 (Two-player Zero-sum Games and Synthetic Data)
April 8, 2025

Adam Smith (based on materials developed with Jonathan Ullman)

Problems with marked with an asterisk **(*)** *are more challenging or open-ended.*

1. In this question we'll explore an alternative strategy where Colin and Rowena try to play best responses to the *previous* action taken by the other player. Consider the following game from the last class' exercises:

$$\begin{bmatrix} +2 & -1 \\ -2 & +3 \end{bmatrix} \tag{1}$$

   (a) If you didn't get to this in the last class, compute equilibrium strategies for Rowena and Colin.

   (b) Suppose Rowena and Colin play the game for T iterations, and for each t, Rowena and Colin both play a best response to the strategy the other player used on the *previous* iteration $t-1$. That is

$$\mathbf{r}_t = \arg\max_{\mathbf{r}} \mathbf{r}^\top M \mathbf{c}_{t-1} \tag{2}$$

$$\mathbf{c}_t = \arg\min_{\mathbf{c}} \mathbf{r}_{t-1}^\top M \mathbf{c} \tag{3}$$

   For concreteness, assume $\mathbf{r}_1 = \mathbf{c}_1 = (1, 0)$, so Rowena starts by playing the top row and Colin starts by playing the left column. Do the sequences of strategies converge to an equilibrium of the game?

   (c) Consider the same setup as above, but now consider the *average* strategies

$$\hat{\mathbf{r}}_T = \frac{1}{T}\sum_{t=1}^{T} \mathbf{r}_t \text{ and } \hat{\mathbf{c}}_T = \frac{1}{T}\sum_{t=1}^{T} \mathbf{c}_t$$

   Do $\hat{\mathbf{r}}_T$ and $\hat{\mathbf{c}}_T$ converge to equilibrium strategies?

2. Prove that if $\mathbf{r}_1, \ldots, \mathbf{r}_T$ and $\mathbf{c}_1, \ldots, \mathbf{c}_T$ are two sequences that each have regret at most α to one another, then $\hat{\mathbf{r}} = \frac{1}{T}\sum_t \mathbf{r}_t$ and $\hat{\mathbf{c}} = \frac{1}{T}\sum_t \mathbf{c}_t$ are 2α-approximate equilibrium strategies.

3. Specialize DualQuery to the case of threshold queries over the universe $\{1, \ldots, D\}$. What do the multiplicative-weights updates for the query-player look like? What do the best-response problems for the data-player look like?

(See over.)

4. The privacy analysis (and hence the accuracy analysis) of the DualQuery is somewhat subtle, and this question will walk you through it at a high level. Recall that in DualQuery, given the query player's strategy $\mathbf{r}_t$, we sample queries $i_{t,1}, \ldots, i_{t,S}$ independently according to $\mathbf{r}_t$ and set $\tilde{\mathbf{r}}_t$ to be the uniform distribution over those queries, and we want to understand how this sampling step ensures privacy without dramatically reducing accuracy.

   (a) First, the data player solves the optimization problem

   $$z_t = \arg\min_z \mathop{\mathbb{E}}_{i \sim \tilde{\mathbf{r}}_t} (-\varphi_i(z))$$

   when in fact they want to optimize with the query-player's actual strategy $\mathbf{r}_t$. Show that if we draw S samples then with high probability, the data player's output will satisfy

   $$z_t \leq \min_z \mathop{\mathbb{E}}_{i \sim \mathbf{r}_t} (-\varphi_i(z)) + \alpha$$

   for some $\alpha = O(\sqrt{\log |\mathcal{U}|/S})$. Thus, the data player's exact best response to $\tilde{\mathbf{r}}_t$ is also an approximate best-response to $\mathbf{r}_t$. [***Hint:*** Chernoff Bounds!]

   (b) The previous statement gives us some guidance on how many samples we need to draw from $\mathbf{r}_t$ in order to get good convergence properties, but why would those samples be private. Argue that if we fix any sequence of responses $z_1, \ldots, z_t$, and the query player computes $\mathbf{r}_t$ using multiplicative weights on the corresponding losses, with update parameter η, then each sample from the distribution $\mathbf{r}_t$ is private for some ε_0 that may depend on η, n, T. [***Hint:*** A few lectures ago we asked a question about the specific form of the distribution maintained by multiplicative weights. Does it look like any type of distribution we've studied in differential privacy?]

   (c) Suppose we run DualQuery for T iterations, with an update step size η, and in each iteration we take S samples from $\mathbf{r}_t$. What bound do we get on the accuracy of the output? [***Hint:*** To get intuition, it's helpful to imagine that the bound you prove in part (a) holds with probability 1, rather than "with high probability", to avoid having to deal with small probabilities of failure in different steps of the algorithm]

   (d) Suppose we set T, η, S in such a way that the resulting algorithm satisfies $(\varepsilon, 0)$-differential privacy. How big do we need to set n in order to guarantee error α. Your bound should have a form like

   $$n \gtrsim \frac{(\log |\mathcal{U}|)^a (\log k)^b}{\varepsilon^c \alpha^d}$$

   which is similar to what we obtain for MWEM up to the specific polynomial factors.

   (e) How would your answer change if we aim for ε, δ privacy (ignoring terms that are polynomial in $\log(1/\delta)$, for simplicity).