# CSCI 678: Theoretical Machine Learning
# Homework 4

**Fall 2024, Instructor: Haipeng Luo**

*This homework is due on* **12/01, 11:59pm**. *See course website for more instructions on finishing and submitting your homework as well as the late policy. Total points: 50*

1. (**Stochastic MAB**) In Lecture 9, we proved that UCB achieves $\overline{\text{Reg}}_n = \mathcal{O}(\sum_{a:\Delta_a>0} \frac{\ln n}{\Delta_a})$, which increases as the gaps decrease and make it harder to distinguish the optimal actions from the rest. However, if an action really has a tiny suboptimality gap, then by definition selecting it does not lead to large regret and thus there is really no point in distinguishing it from the optimal actions. Building on this intuition, in this problem you will further prove that UCB indeed also guarantees $\mathcal{O}(\sqrt{nK\ln n})$ pseudo regret at the same time, regardless of the values of the gaps.

   (a) (5pts) Recall (from Theorem 1 of Lecture 9) that with probability at least $1 - \frac{2K}{n}$, we have $m_n(a) \le \frac{16\ln n}{\Delta_a^2} + 1$ for every action $a$ where $m_n(a)$ is the total number of times action $a$ is pulled. Use this fact to prove the following pseudo regret bound

$$\overline{\text{Reg}}_n \le 3K + \Delta n + \sum_{a:\Delta_a>\Delta} \frac{16\ln n}{\Delta_a}, \tag{1}$$

   where $\Delta \in [0,1]$ is an arbitrary threshold. (This matches the earlier intuition that we do not care about distinguishing an action with a small gap from the optimal action.)

   *Proof.* Similar to the proof of Theorem 1, we denote by $E$ the high probability event under which $m_n(a) \le \frac{16\ln n}{\Delta_a^2} + 1$ for all $a$, and bound the pseudo regret as

$$\begin{aligned}\overline{\text{Reg}}_n &\le \Pr(E) \times \sum_{a:\Delta_a>0} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right] + \Pr(\neg E) \times n \\ &\le \sum_{a:\Delta_a>0} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right] + 2K.\end{aligned}$$

   To continue, we separate the term $\sum_{a:\Delta_a>0} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right]$ into two parts:

$$\sum_{a:\Delta_a>0} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right] = \sum_{a:\Delta_a\le\Delta} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right] + \sum_{a:\Delta_a>\Delta} \Delta_a \mathbb{E}\left[m_n(a) \mid E\right].$$

   Next, we bound the first term trivially by $\Delta \sum_{a:\Delta_a\le\Delta} \mathbb{E}\left[m_n(a) \mid E\right] \le \Delta n$ (since $\sum_{a=1}^K m_n(a) = n$), and the second term by

$$\sum_{a:\Delta_a>\Delta} \Delta_a \times \left(\frac{16\ln n}{\Delta_a^2} + 1\right) \le \sum_{a:\Delta_a>\Delta} \frac{16\ln n}{\Delta_a} + K.$$

   Combining everything proves the claim. ☐

(b) (2pts) Pick an appropriate value of $\Delta$ and conclude the following bound

$$\overline{\text{Reg}}_n = \mathcal{O}(\sqrt{nK \ln n}).$$

(You can assume $n \geq K$.)

*Proof.* First, note that the bound from the last question can be further bounded as

$$\overline{\text{Reg}}_n \leq 3K + \Delta n + \frac{16K \ln n}{\Delta}.$$

Therefore, picking $\Delta = \min\left\{1, \sqrt{\frac{K \ln n}{n}}\right\}$ (which is order-optimal) then proves $\overline{\text{Reg}}_n = \mathcal{O}(\sqrt{nK \ln n})$. □

2. (**Multiclass Perceptron**) In this exercise, you need to analyze variants of the Perceptron algorithm for *multiclass* classification, with either full information or bandit information. Specifically, consider a sequence of examples $x_1, \ldots, x_n \in B_2^d$ with labels $y_1, \ldots, y_n \in [K]$ where $K$ is the number of possible classes. We assume that the following multiclass margin assumption holds: there exists a constant $\gamma > 0$ and $K$ weight vectors $\theta_\star^1, \ldots, \theta_\star^K \in B_2^d$ such that for each $t = 1, \ldots, n$:

$$\left\langle \theta_\star^{y_t}, x_t \right\rangle \geq \left\langle \theta_\star^k, x_t \right\rangle + \gamma, \quad \forall k \neq y_t.$$

In other words, the predictor $\mathrm{argmax}_k \left\langle \theta^k, x_t \right\rangle$ makes perfect predictions for this dataset with $\gamma$ margin. Now, consider the following learning protocol:

For $t = 1, \ldots, n$:
- receive $x_t$ and predict $s_t \in [K]$;
- observe $\begin{cases} y_t & \text{in the full-information setting} \\ \mathbf{1}\{s_t \neq y_t\} \text{ (i.e., if the prediction is correct)} & \text{in the bandit setting} \end{cases}$

In either case, we care about the total number of mistakes $M = \sum_{t=1}^n \mathbf{1}\{s_t \neq y_t\}$.

(a) In the full information setting, one can apply the following multiclass Perceptron algorithm, a natural generalization of its binary version studied in Lecture 7. Note that when the algorithm predicts correctly, the last update step in fact does nothing (similarly to the binary version).

**Algorithm 1:** Multiclass Perceptron

Initialize $\theta^1 = \cdots = \theta^K = \mathbf{0}$.
For $t = 1, \ldots, n$:
- receive $x_t$ and find $k_t \in \mathrm{argmax}_{k \in [K]} \left\langle \theta^k, x_t \right\rangle$;
- predict $s_t = k_t$;
- receive $y_t$ and update

$$\theta^{y_t} \leftarrow \theta^{y_t} + x_t \quad \text{and} \quad \theta^{k_t} \leftarrow \theta^{k_t} - x_t.$$

Follow the steps below to prove $M \leq \frac{2K}{\gamma^2}$ for this algorithm.

i. (6pts) Similar to the binary case, we need to analyze the evolution of the quantities $\sum_{k=1}^K \left\langle \theta^k, \theta_\star^k \right\rangle$ and $\sum_{k=1}^K \left\|\theta^k\right\|_2^2$. To this end, denote the value of the weight vectors $\theta^1, \ldots, \theta^K$ at the beginning of round $t$ by $\theta_t^1, \ldots, \theta_t^K$. Under the margin assumption, prove the following two facts for any $t = 1, \ldots, n$:

$$\sum_{k=1}^K \left\langle \theta_{t+1}^k, \theta_\star^k \right\rangle \geq \sum_{k=1}^K \left\langle \theta_t^k, \theta_\star^k \right\rangle + \gamma \mathbf{1}\{s_t \neq y_t\},$$

and

$$\sum_{k=1}^K \left\|\theta_{t+1}^k\right\|_2^2 \leq \sum_{k=1}^K \left\|\theta_t^k\right\|_2^2 + 2\mathbf{1}\{s_t \neq y_t\}.$$

*Proof.* If $y_t = s_t$, then clearly $\theta_{t+1}^k = \theta_t^k$ for all $k$ and both statements become trivial. Now suppose $y_t \neq s_t = k_t$, so that $\theta_{t+1}^{y_t} = \theta_t^{y_t} + x_t$, $\theta_{t+1}^{k_t} = \theta_t^{k_t} - x_t$, and $\theta_{t+1}^k = \theta_t^k$ for $k \notin \{y_t, k_t\}$. In this case, we have

$$\sum_{k=1}^K \left\langle \theta_{t+1}^k, \theta_\star^k \right\rangle = \sum_{k=1}^K \left\langle \theta_t^k, \theta_\star^k \right\rangle + \left\langle x_t, \theta_\star^{y_t} - \theta_\star^{k_t} \right\rangle \geq \sum_{k=1}^K \left\langle \theta_t^k, \theta_\star^k \right\rangle + \gamma,$$

where the last step is by the margin assumption and $k_t \neq y_t$. This proves the first statement. The second statement holds since

$$\sum_{k=1}^K \left\|\theta_{t+1}^k\right\|_2^2 = \sum_{k=1}^K \left\|\theta_t^k\right\|_2^2 + 2\left\langle \theta_t^{y_t} - \theta_t^{k_t}, x_t \right\rangle + 2\left\|x_t\right\|_2^2 \leq \sum_{k=1}^K \left\|\theta_t^k\right\|_2^2 + 2$$

where the last steps uses the fact $\|x_t\|_2 \leq 1$ and the definition of $k_t$. □

ii. (3pts) Combine the two facts in the last question to conclude $M \leq \frac{2K}{\gamma^2}$ (Hint: you will need to use the Cauchy-Schwarz inequality.)

*Proof.* Repeatedly applying the two facts from the last question shows $\gamma M \leq \sum_{k=1}^{K} \langle \theta_{n+1}^k, \theta_\star^k \rangle$ and $\sum_{k=1}^{K} \|\theta_{n+1}^k\|_2^2 \leq 2M$. To connect these two, we apply the Cauchy-Schwarz inequality together with the fact $\|\theta_\star^k\| \leq 1$ to show

$$\sum_{k=1}^{K} \langle \theta_{n+1}^k, \theta_\star^k \rangle \leq \sqrt{\sum_{k=1}^{K} \|\theta_{n+1}^k\|_2^2} \sqrt{\sum_{k=1}^{K} \|\theta_\star^k\|_2^2} \leq \sqrt{K \sum_{k=1}^{K} \|\theta_{n+1}^k\|_2^2}.$$

Combining everything shows $\gamma M \leq \sqrt{2KM}$, which implies $M \leq \frac{2K}{\gamma^2}$ as claimed. □

(b) In the bandit setting, we make the following two changes to Algorithm 1: 1) first, in light of the exploration versus exploitation trade-off, it is natural to randomize the algorithm and explore every label with at least some small probability $\alpha$; 2) second, the update $\theta^{y_t} \leftarrow \theta^{y_t} + x_t$ becomes invalid if the prediction is incorrect (since we do not know what $y_t$ is), so we only do this update when we predict correctly, and we scale the update with the inverse probability of selecting the correct label, just like the idea of importance-weighted estimator in Exp3. The final algorithm is shown below.

**Algorithm 2:** Multiclass Perceptron with Bandit Feedback

Input: exploration parameter $\alpha \in (0, \frac{1}{2K}]$.
Initialize $\theta^1 = \cdots = \theta^K = \mathbf{0}$.
For $t = 1, \ldots, n$:

- receive $x_t$ and find $k_t \in \operatorname{argmax}_{k \in [K]} \langle \theta^k, x_t \rangle$;
- predict $s_t$ drawn from $p_t$, where $p_t(k) = (1 - \alpha K)\mathbf{1}\{k = k_t\} + \alpha, \forall k$;
- receive $\mathbf{1}\{s_t \neq y_t\}$ and update

$$\theta^{y_t} \leftarrow \theta^{y_t} + \frac{x_t \mathbf{1}\{s_t = y_t\}}{p_t(y_t)} \quad \text{and} \quad \theta^{k_t} \leftarrow \theta^{k_t} - x_t.$$

Follow the steps below to prove that this algorithm makes at most $\mathcal{O}\left(\frac{K\sqrt{n}}{\gamma^2}\right)$ mistakes in expectation. We will again use the notation $\theta_t^1, \ldots, \theta_t^K$ to denote the value of the weight vectors $\theta^1, \ldots, \theta^K$ at the beginning of round $t$, and study the evolution of the quantities $\sum_{k=1}^{K} \langle \theta^k, \theta_\star^k \rangle$ and $\sum_{k=1}^{K} \|\theta^k\|_2^2$ (*in expectation* this time).

i. (3pts) Under the margin assumption, prove the following for any $t = 1, \ldots, n$:

$$\mathbb{E}\left[\sum_{k=1}^{K} \langle \theta_{t+1}^k, \theta_\star^k \rangle\right] \geq \mathbb{E}\left[\sum_{k=1}^{K} \langle \theta_t^k, \theta_\star^k \rangle\right] + \gamma \mathbb{E}\left[\mathbf{1}\{k_t \neq y_t\}\right].$$

*Proof.* Condition on the history before round $t$ so that $\theta_t^1, \ldots, \theta_t^K$ and $k_t$ are fixed, and let $\mathbb{E}_t$ be the conditional expectation. Then we have

$$\begin{aligned}
\mathbb{E}_t\left[\sum_{k=1}^{K} \langle \theta_{t+1}^k, \theta_\star^k \rangle\right] &= \left[\sum_{k=1}^{K} \langle \theta_t^k, \theta_\star^k \rangle\right] + \mathbb{E}_t\left[\left\langle \frac{x_t \mathbf{1}\{s_t = y_t\}}{p_t(y_t)}, \theta_\star^{y_t} \right\rangle\right] - \langle x_t, \theta_\star^{k_t} \rangle \\
&= \left[\sum_{k=1}^{K} \langle \theta_t^k, \theta_\star^k \rangle\right] + \langle x_t, \theta_\star^{y_t} - \theta_\star^{k_t} \rangle \\
&\geq \left[\sum_{k=1}^{K} \langle \theta_t^k, \theta_\star^k \rangle\right] + \mathbf{1}\{k_t \neq y_t\}\gamma
\end{aligned}$$

where the last step uses the margin assumption. Further taking expectation over the history finishes the proof. □

ii. (8pts) Next, prove the following for any $t = 1, \ldots, n$:

$$\mathbb{E}\left[\sum_{k=1}^{K} \left\|\theta_{t+1}^k\right\|_2^2\right] \le \mathbb{E}\left[\sum_{k=1}^{K} \left\|\theta_t^k\right\|_2^2\right] + \frac{\mathbb{E}\left[\mathbf{1}\left\{k_t \ne y_t\right\}\right]}{\alpha} + 1.$$

Hint: consider the two cases $k_t \ne y_t$ and $k_t = y_t$ separately.

*Proof.* Again, condition on the history before round $t$ so that $\theta_t^1, \ldots, \theta_t^K$ and $k_t$ are fixed, and let $\mathbb{E}_t$ be the conditional expectation. If $k_t \ne y_t$, then

$$\begin{aligned}
&\mathbb{E}_t\left[\sum_{k=1}^{K} \left\|\theta_{t+1}^k\right\|_2^2\right] - \sum_{k=1}^{K} \left\|\theta_t^k\right\|_2^2 \\
&= \mathbb{E}_t\left[\left\|\theta_t^{y_t} + \frac{x_t \mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)}\right\|_2^2\right] - \left\|\theta_t^{y_t}\right\|_2^2 + \left\|\theta_t^{k_t} - x_t\right\|_2^2 - \left\|\theta_t^{k_t}\right\|_2^2 \\
&= 2\mathbb{E}_t\left[\left\langle x_t, \frac{\mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)}\theta_t^{y_t} - \theta_t^{k_t}\right\rangle\right] + \mathbb{E}_t\left[\frac{\mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)^2}\right]\left\|x_t\right\|_2^2 + \left\|x_t\right\|_2^2 \\
&= 2\left\langle x_t, \theta_t^{y_t} - \theta_t^{k_t}\right\rangle + \left(\frac{1}{p_t(y_t)} + 1\right)\left\|x_t\right\|_2^2 \\
&\le \frac{1}{p_t(y_t)} + 1 && \text{(By definition of } k_t \text{ and } \|x_t\| \le 1\text{)} \\
&= \frac{1}{\alpha} + 1. && (p_t(k) = \alpha \text{ for } k \ne k_t)
\end{aligned}$$

On the other hand, if $k_t = y_t$, then

$$\begin{aligned}
&\mathbb{E}_t\left[\sum_{k=1}^{K} \left\|\theta_{t+1}^k\right\|_2^2\right] - \sum_{k=1}^{K} \left\|\theta_t^k\right\|_2^2 \\
&= \mathbb{E}_t\left[\left\|\theta_t^{y_t} + \left(\frac{\mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)} - 1\right)x_t\right\|_2^2\right] - \left\|\theta_t^{y_t}\right\|_2^2 \\
&= 2\mathbb{E}_t\left[\left(\frac{\mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)} - 1\right)\left\langle x_t, \theta_t^{y_t}\right\rangle\right] + \mathbb{E}_t\left[\left(\frac{\mathbf{1}\left\{s_t = y_t\right\}}{p_t(y_t)} - 1\right)^2\right]\left\|x_t\right\|_2^2.
\end{aligned}$$

The first conditional expectation above is simply 0, while the second one is

$$p_t(y_t) \times \left(\frac{1}{p_t(y_t)} - 1\right)^2 + (1 - p_t(y_t)) = \frac{1}{p_t(y_t)} - 1.$$

Therefore, using the facts $p_t(y_t) = p_t(k_t) \ge 1 - \alpha K$ and $\|x_t\|_2 \le 1$ we arrive at

$$\mathbb{E}_t\left[\sum_{k=1}^{K} \left\|\theta_{t+1}^k\right\|_2^2\right] - \sum_{k=1}^{K} \left\|\theta_t^k\right\|_2^2 \le \frac{1}{1 - \alpha K} - 1 \le 1,$$

where the last step is due to $\alpha \le \frac{1}{2K}$. Combining the two cases we have

$$\mathbb{E}_t\left[\sum_{k=1}^{K} \left\|\theta_{t+1}^k\right\|_2^2\right] - \sum_{k=1}^{K} \left\|\theta_t^k\right\|_2^2 \le \left(\frac{1}{\alpha} + 1\right)\mathbf{1}\left\{k_t \ne y_t\right\} + \mathbf{1}\left\{k_t = y_t\right\} = \frac{\mathbf{1}\left\{k_t \ne y_t\right\}}{\alpha} + 1.$$

Further taking expectation over the history finishes the proof. □

iii. (5pts) Combine the results from the last two questions to show $\gamma\mathbb{E}[N] \le \sqrt{K\left(\frac{\mathbb{E}[N]}{\alpha} + n\right)}$ where $N = \sum_{t=1}^{n} \mathbf{1}\left\{k_t \ne y_t\right\}$. Further solve for $\mathbb{E}[N]$ to show $\mathbb{E}[N] \le \frac{K}{\alpha\gamma^2} + \frac{\sqrt{Kn}}{\gamma}$.

*Proof.* Repeatedly applying the result from the first question shows $\gamma\mathbb{E}[N] \leq \mathbb{E}\left[\sum_{k=1}^{K}\left\langle \theta_{n+1}^k, \theta_\star^k \right\rangle\right]$, which is further bounded by

$$\mathbb{E}\left[\sqrt{\sum_{k=1}^{K}\left\|\theta_{n+1}^k\right\|_2^2}\right]\sqrt{\sum_{k=1}^{K}\|\theta_\star^k\|_2^2} \leq \mathbb{E}\left[\sqrt{K\sum_{k=1}^{K}\left\|\theta_{n+1}^k\right\|_2^2}\right] \leq \sqrt{K\mathbb{E}\left[\sum_{k=1}^{K}\left\|\theta_{n+1}^k\right\|_2^2\right]}$$

where the first step is by Cauchy-Schwarz inequality, the second step is by $\left\|\theta_\star^k\right\|_2 \leq 1$, and the last step is by Jensen's inequality. Now, repeatedly applying the result from the second question shows

$$\gamma\mathbb{E}[N] \leq \sqrt{K\mathbb{E}\left[\sum_{t=1}^{n}\frac{\mathbf{1}\{k_t \neq y_t\}}{\alpha} + 1\right]} = \sqrt{K\left(\frac{\mathbb{E}[N]}{\alpha} + n\right)}.$$

Squaring both sides and rearranging leads to the following quadratic inequality in terms of $\mathbb{E}[N]$:

$$\gamma^2\mathbb{E}[N]^2 - \frac{K}{\alpha}\mathbb{E}[N] - Kn \leq 0.$$

Solving for it gives

$$\begin{aligned}\mathbb{E}[N] &\leq \frac{1}{2\gamma^2}\left(\frac{K}{\alpha} + \sqrt{\frac{K^2}{\alpha^2} + 4\gamma^2 Kn}\right) \\ &\leq \frac{1}{2\gamma^2}\left(\frac{K}{\alpha} + \sqrt{\frac{K^2}{\alpha^2}} + \sqrt{4\gamma^2 Kn}\right) = \frac{K}{\alpha\gamma^2} + \frac{\sqrt{Kn}}{\gamma}.\end{aligned}$$

This finishes the proof. □

iv. (4pts) Finally, use the result from the last step to prove $\mathbb{E}[M] \leq \frac{K}{\alpha\gamma^2} + \frac{\sqrt{Kn}}{\gamma} + \alpha nK$, and pick an appropriate value of $\alpha$ to conclude $\mathbb{E}[M] = \mathcal{O}\left(\frac{K\sqrt{n}}{\gamma} + \frac{K^2}{\gamma^2}\right)$.

*Proof.* It suffices the connect $N$ and $M$ as follows:

$$\begin{aligned}\mathbb{E}[M] = \mathbb{E}\left[\sum_{t=1}^{n}\mathbf{1}\{s_t \neq y_t\}\right] &= \mathbb{E}\left[\sum_{t=1}^{n}\sum_{k=1}^{K}p_t(k)\mathbf{1}\{k \neq y_t\}\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{n}\left(p_t(k_t)\mathbf{1}\{k_t \neq y_t\} + \alpha\sum_{k\neq k_t}\mathbf{1}\{k \neq y_t\}\right)\right] \\ &\leq \mathbb{E}[N] + \alpha nK \leq \frac{K}{\alpha\gamma^2} + \frac{\sqrt{Kn}}{\gamma} + \alpha nK.\end{aligned}$$

Picking the optimal choice of $\alpha = \min\left\{\frac{1}{\gamma\sqrt{n}}, \frac{1}{2K}\right\}$ proves the claimed mistake bound. □

3. (**Partial Monitoring**) Recall the dynamic pricing problem discussed in Lecture 9 and consider a simplified case with only 3 possible prices ($1, 2$, or 3 dollars). The loss matrix and feedback matrix are thus

$$\ell = \begin{pmatrix} 0 & 1 & 2 \\ c & 0 & 1 \\ c & c & 0 \end{pmatrix} \qquad \text{and} \qquad \Phi = \begin{pmatrix} ✓ & ✓ & ✓ \\ ✗ & ✓ & ✓ \\ ✗ & ✗ & ✓ \end{pmatrix}$$

for some storage cost $c > 0$. The cell decomposition of this problem is illustrated in the following picture, where we show the simplex $\Delta(3)$ by considering only the first two coordinates $u_1$ and $u_2$. Clearly, all 3 actions are Pareto-optimal, and every two actions are neighbors.

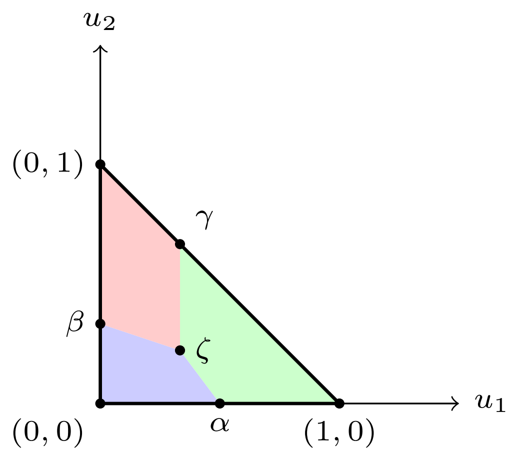


(a) (3pts) State which colored region in the cell decomposition picture corresponds to cell $C_1$, $C_2$, and $C_3$ respectively. Briefly explain why.

*Proof.* The green region contains $u = (1, 0, 0)$, under which action 1 is clearly the optimal action, meaning that this region must be $C_1$. Similarly, the red region contains $u = (0, 1, 0)$, under which action 2 is the optimal action, so this region is $C_2$, and the blue region must be $C_3$. □

(b) (4pts) Calculate the coordinates of the four points $\alpha$, $\beta$, $\gamma$, and $\zeta$ shown on the cell decomposition picture.

*Proof.* Let $(x, 0, 1-x) \in \Delta(3)$ be the distribution that $\alpha$ represents. The fact that it sits on $C_1 \cap C_3$ implies $\langle \ell_1, (x, 0, 1-x)\rangle = \langle \ell_3, (x, 0, 1-x)\rangle$. Plugging in $\ell_1 = (0, 1, 2)$ and $\ell_3 = (c, c, 0)$ gives $2(1-x) = xc$. Solving for $x$ gives $x = \frac{2}{c+2}$ and thus $\alpha = (\frac{2}{c+2}, 0)$.
Similarly, let $(0, y, 1-y)$ be the distribution that $\beta$ represents. It sits on $C_2 \cap C_3$ and thus $\langle \ell_2, (0, y, 1-y)\rangle = \langle \ell_3, (0, y, 1-y)\rangle$. Solving for $y$ gives $\beta = (0, \frac{1}{c+1})$.
Next, let $(z, 1-z, 0)$ be the distribution that $\gamma$ represents. It sits on $C_1 \cap C_2$ and thus $\langle \ell_1, (z, 1-z, 0)\rangle = \langle \ell_2, (z, 1-z, 0)\rangle$. Solving for $z$ gives $\gamma = (\frac{1}{c+1}, \frac{c}{c+1})$.
Finally, to find the coordinates of $\zeta$, we find the distribution $u \in \Delta(3)$ such that $\langle \ell_1, u\rangle = \langle \ell_2, u\rangle = \langle \ell_3, u\rangle$, which leads to $u = (\frac{1}{c+1}, \frac{c}{(c+1)^2}, \frac{c^2}{(c+1)^2})$, and thus $\zeta = (\frac{1}{c+1}, \frac{c}{(c+1)^2})$. □

(c) (4pts) Prove that the following two action pairs are both locally observable: 1 and 2, 2 and 3.

*Proof.* The signal matrices are

$$S_1 = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \quad S_2 = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \end{pmatrix}, \quad S_3 = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$$

(note that swapping the two rows is also correct). Therefore, we have

$$\ell_1 - \ell_2 = (-c, 1, 1) = (1, -c)S_2$$

and

$$\ell_2 - \ell_3 = (0, -c, 1) = (0, c, 1, -c)\begin{pmatrix} S_2 \\ S_3 \end{pmatrix},$$

which by definition implies their local observability. □

(d) (3pts) The results from the last question imply that actions 1 and 3 must be globally observable. Now, prove that they are not locally observable. (This implies that this is a globally observable but not locally observable partial monitoring instance.)

*Proof.* Note that $\ell_1 - \ell_3 = (-c, 1 - c, 2)$, with the first two coordinates being distinct. On the other hand, the rows of $S_1$ and $S_3$ all have identical values for the first two coordinates. This means that $\ell_1 - \ell_3$ cannot be in the row space of $S_{N_{13}}$ and thus actions 1 and 3 are not locally observable. □