Mehryar Mohri
Advanced Machine Learning 2025
Courant Institute of Mathematical Sciences
Homework assignment 1
March 04, 2025
Due: March 18, 2025

A RWM bound

Using the same proof as the one given in class, prove that $R_T \leq 4\sqrt{\mathcal{L}_T^{\min} \log N}$ for an appropriate choice of the parameter β.

B Resource allocation games

Two competing AI algorithms (Player A and Player B) are tasked with managing a shared pool of computational resources. These resources are essential for both algorithms to perform their respective tasks, which contribute to a larger, overarching system.

Each algorithm can choose to either: "Prioritize Efficiency" (PE): Focus on optimizing its own resource usage, potentially leaving less for the other; or "Promote Sharing" (PS): Distribute resources more evenly, potentially sacrificing some of its own immediate efficiency. The system's overall performance depends on a balance between individual efficiency and collaborative resource sharing.

If both algorithms prioritize efficiency (PE, PE), they might achieve high individual performance, but the overall system could become unstable due to resource contention, resulting in a moderate negative payoff for both. If both algorithms promote sharing (PS, PS), the system remains stable, but individual performance might be slightly lower, resulting in a moderate positive payoff for both. If one algorithm prioritizes efficiency while the other promotes sharing (PE, PS or PS, PE), the efficient algorithm gains a significant advantage, while the sharing algorithm suffers a loss. The overall system may become unstable.

The payoff matrix of the game is the following:

Player A/B	PE	PS
PE	(-2, -2)	(+3, 0))
PS	(0, +3)	(+2, +2)

1. Nash equilibria: show that (PE,PS) and (PS,PE) are Nash equilibria. Give an interpretation of these equilibria. Show that there is no pure Nash equilibrium where both players choose the same strategy.

2. Correlated equilibria: Show that a correlated equilibria can be achieved if a central resource manager randomly instructs the algorithms to alternate between prioritizing efficiency and promoting sharing. This would ensure system stability and fair resource distribution over time. This would look like a 50% chance of (PE,PS) and a 50% chance of (PS,PE). Show that this would create a much better average payoff for both players.

3. We now consider a more advanced resource allocation game where the two algorithms have more granular control over their resource allocation. Their strategies or actions are:

   - Aggressive Acquisition (AA): Prioritize maximum resource capture, risking severe contention.

- Balanced Optimization (BO): Seek efficiency while maintaining a moderate level of resource sharing.
- Collaborative Sharing (CS): Focus on equitable resource distribution, even at the cost of significant individual performance.

The payoff matrix is

Player A/B	AA	BO	CS
AA	(-5, -5)	(+4, -2)	(+2, -4)
BO	(-2, 4)	(+1, +1)	(+3, 0)
CS	(-4, +2)	(0, +3)	(0, 0)

(a) Give the pure Nash equilibria of the game. Bonus question: give all mixed Nash equilibria.

(b) A central coordinator (or a shared signal) could recommend the following distribution: (BO, BO): 1/3 probability, (AA, BO): 1/3 probability, and (BO, AA): 1/3 probability. Show that this is a correlated equilibrium. What are its main benefits?

C Mirror descent

Consider the function $\Phi\colon C \to \mathbb{R}$ defined over $C = (\mathbb{R}_+^*)^d$ by:

$$\Phi(x) = \sum_{i=1}^{d} x_i \log(x_i) + \lambda \sum_{i=1}^{d} |x_i|,$$

where λ is a non-negative constant. Show that Φ can be used to define a mirror descent algorithm over the simplex, assuming that the loss functions f_t selected by the adversary are convex and Lipschitz with respect to ℓ_∞-norm (give full justifications). Give the explicit expression of the Bregman divergence and MD updates. Indicate the regret guarantees of the algorithm.