MIT 6.S890 — Topics in Multiagent Learning Tue, Nov 12th 2024

Lecture 15

Computation of exact equilibria

Instructor: Prof. Gabriele Farina (✉ gfarina@mit.edu)*

Today we briefly discuss an aspect that we have not investigated directly. We have talked many times about how learning enables us to compute equilibria in games. In particular, we have connected coarse correlated equilibrium (CCE) approximation to the ratio between the regret of the learning algorithm and time. In practice, this means that most learning procedures recover a CCE at a rate of $O(1/\sqrt{T})$ or, in the best case, $O(1/T)$ using optimistic dynamics (Lecture 5). So, while learning methods are today the most scalable way of computing CCEs in large games, it is clear that they are only suitable for *low-precision* computations, *i.e.*, for relatively small values of ϵ. This begs the following question: are there algorithms that can compute *high-precision* CCEs, and other types of equilibria, in games?

1 The main ideas

The answer to the above question is positive, and comes from a very neat construction by Papadimitriou, C. H., & Roughgarden, T. [PR08]. Long story short, for any sufficiently nice choice of convex strategy set $\mathcal{X}_i$ per player, we know how to construct an algorithm that computes an ϵ-equilibrium in a number of steps that scales with the *logarithm* of $1/\epsilon$.

1.1 Hart & Schmeidler's proof of existence of (coarse) correlated equilibria

So far in the course, we have proven the existence of coarse correlated and correlated equilibria through the existence of Nash equilibria. However, one might wonder if there is a more direct way to prove the existence of CEs and CCEs, which does not rely on the existence of a much harder notion. The answer is yes, and the idea comes from a very neat proof by Hart, S., & Schmeidler, D. [HS89], which includes some ideas that will set the stage for the Ellipsoid-Against-Hope algorithm. The proof is based on linear programming duality.

While the original proof of Hart, S., & Schmeidler, D. [HS89] is for CE, I will present here a version of the proof simplified for the case of CCEs.

As a reminder, by definition a coarse correlated equilibrium is a distribution $\mu \in \Delta(S_1 \times \ldots \times S_n)$ such that

$$\mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i})] \le \mathbb{E}_{s\sim\mu}[u_i(s_i, s_{-i})] \qquad \forall i \in [n], s_i' \in S_i.$$

Finding a CCE is then equivalent to finding a distribution μ such that

$$\min_{\mu} \max_{\substack{i\in[n]\\ s_i'\in S_i}} \mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \le 0.$$

*These notes are class material that has not undergone formal peer review. The TA and I are grateful for any reports of typos.

How can we prove the existence of such a μ without resorting to the existence of Nash equilibria? The rescue comes from the minimax theorem (which itself can be shown using linear programming duality or, as we discussed in Lecture 4, using the existence of no-regret algorithms).

Before we can use the minimax theorem, we have to "convexify" the inner problem however, since the maximum is currently on a discrete set. To convexity the problem, we will simply allow the possibility for the internal maximumization problem to propose a *distribution* ν over deviations (i, s_i), and we will rewrite the problem as

$$\min_{\mu} \max_{\nu} \mathbb{E}_{s\sim\mu}\mathbb{E}_{(i,s_i')\sim\nu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \leq 0.$$

By the minimax theorem and swapping the order of the expectations, the above min-max value is equal to

$$\max_{\nu} \min_{\mu} \mathbb{E}_{(i,s_i')\sim\nu}\mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})].$$

Can we show that this value is ≤ 0? The answer is yes, and constructive.

Theorem 1.1. The following inequality holds:

$$\max_{\nu} \min_{\mu} \mathbb{E}_{(i,s_i')\sim\nu}\mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \leq 0.$$

Proof. The key here is to pick μ in a way that depends on ν. In particular, we will pick μ to be the *product* distribution that outputs

$$(s_1, ..., s_n) \text{ with probability} \propto \nu_{1,s_1} \cdot ... \cdot \nu_{n,s_n}.$$

With this choice and some simple manipulations,

$$\begin{aligned}
&\mathbb{E}_{(i,s_i')\sim\nu}\mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \\
&\qquad = \sum_i \sum_{s_i' \in S_i} \nu_{i,s_i'} \mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \\
&\qquad = \sum_i \sum_{s_i' \in S_i} \nu_{i,s_i'} \left(\mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i})] - \mathbb{E}_{s\sim\mu}[u_i(s_i, s_{-i})] \right) \\
&\qquad = \sum_i \left(\sum_{s_i' \in S_i} \nu_{i,s_i'} \mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i})] - \sum_{s_i' \in S_i} \nu_{i,s_i'} \mathbb{E}_{s\sim\mu}[u_i(s_i, s_{-i})] \right) \\
&\qquad = \sum_i \left(\underbrace{\sum_{s_i' \in S_i} \nu_{i,s_i'} \mathbb{E}_{s\sim\mu}[u_i(s_i', s_{-i})]}_{(\spadesuit)} - \underbrace{\left(\sum_{s_i' \in S_i} \nu_{i,s_i'} \right) \mathbb{E}_{s\sim\mu}[u_i(s_i, s_{-i})]}_{(\clubsuit)} \right)
\end{aligned}$$

We now expand ($\spadesuit$), using the symbol $\nu_{-i,s_{-i}}$ to mean the product of $\nu_{1,s_1}, ..., \nu_{i-1,s_{i-1}}, \nu_{i+1,s_{i+1}}, ..., \nu_{n,s_n}$.

$$\begin{aligned}
(\spadesuit) &= \sum_{s_i' \in S_i} \nu_{i,s_i'} \sum_{s_i \in S_i} \sum_{s_{-i} \in S_{-i}} \nu_{i,s_i} \nu_{-i,s_{-i}} u_i(s_i', s_{-i}) \\
&= \sum_{s_i' \in S_i} \nu_{i,s_i'} \left(\left(\sum_{s_{-i} \in S_{-i}} \nu_{-i,s_{-i}} u_i(s_i', s_{-i}) \right) \left(\sum_{s_i \in S_i} \nu_{i,s_i} \right) \right) \\
&= \left(\sum_{s_i' \in S_i} \nu_{i,s_i'} \sum_{s_{-i} \in S_{-i}} \nu_{-i,s_{-i}} u_i(s_i', s_{-i}) \right) \left(\sum_{s_i \in S_i} \nu_{i,s_i} \right) \\
&= \left(\sum_{s_i' \in S_i} \sum_{s_{-i} \in S_{-i}} \nu_{i,s_i'} \nu_{-i,s_{-i}} u_i(s_i', s_{-i}) \right) \left(\sum_{s_i \in S_i} \nu_{i,s_i} \right) = (\clubsuit).
\end{aligned}$$

Hence,

$$\max_{\nu} \min_{\mu} \mathbb{E}_{(i,s_i') \sim \nu} \mathbb{E}_{s \sim \mu} [u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] \leq \max_{\nu} 0 = 0.$$

□

1.2 Turning the proof into an algorithm

We now sketch the idea behind the Ellipsoid-Against-Hope algorithm. The full details behind the algorithm are a bit involved, and the machinery a bit technical. However, the important idea is that one can think of the Ellipsoid-Against-Hope algorithm as a computational, constructive proof of the minimax theorem. In particular:

- The proof of Hart and Schmeidler constructively gives a strong response of the mediator to any deviation. In particular, given any ν, there exists a μ that gives the mediator value ≤ 0.
- The ellipsoid method can then be used to find a μ^* such that for all ν the value of the mediator value is ≤ 0.

In more detail, what Theorem 1.1 implies is that the following open polytope must be empty:

$$\left\{ \nu \in \Delta\{(i, s_i') : i \in [n], s_i' \in S_i\} : \mathbb{E}_{s \sim \mu}[u_i(s_i', s_{-i}) - u_i(s_i, s_{-i})] > 0 \quad \forall \mu \in \Delta(S_1 \times \ldots \times S_n) \right\}.$$

Furthermore, for any ν, we know how to prove that at least one of the constraints is violated. The key idea is then to use the ellipsoid method to *certify* the emptiness of the polytope. Normally, the ellipsoid method is used to find a point in a set, but in our case, the point does not exist and we want to use the ellipsoid method to isolate constraints that prove the emptiness of the set. For this reason, the algorithm was called Ellipsoid-Against-Hope by Papadimitriou, C. H., & Roughgarden, T. [PR08].

The ellipsoid will maintain a search space which can be thought of as a suitable subset of the deviator's set. At every iteration, the algorithm will compute the center point ν of the set. Then, it will find a violated constraint using the mediator's strategy in the proof of Theorem 1.1. The violated constraint implies that the deviator set must be curtailed, and the ellipsoid will be updated accordingly reducing the size of the search space by a constant. The algorithm will continue until the search space is small enough to guarantee that the set is empty. In the process, it takes $O(\log(1/\epsilon))$ iterations for the search space to shrink to size ϵ. Along the way, the algorithm will have produced several violated constraints, each of which is associated with a mediator strategy μ. It can be shown using the Farkas lemma, that at least one convex combination of these μ is a CCE.

2 Bibliographic remarks

If you are curious to read more, the following papers contains extensions and refinements of the idea of Ellipsoid-Against-Hope.

[PR08] C. H. Papadimitriou and T. Roughgarden, “Computing correlated equilibria in multi-player games,” *Journal of the ACM (JACM)*, vol. 55, no. 3, pp. 1–29, 2008.

[HS89] S. Hart and D. Schmeidler, “Existence of Correlated Equilibria,” *Mathematics of Operations Research*, vol. 14, no. 1, pp. 18–25, 1989, Accessed: Nov. 07, 2024. [Online]. Available: http://www.jstor.org/stable/3689835

[JL11] A. X. Jiang and K. Leyton-Brown, “Polynomial-time computation of exact correlated equilibrium in compact games,” in *Proceedings of the 12th ACM conference on Electronic commerce*, 2011, pp. 119–126.

[FP24] G. Farina and C. Pipis, “Polynomial-Time Computation of Exact Φ-Equilibria in Polyhedral Games”, in *Neural Information Processing Systems (NeurIPS)*, 2024.

[HS08] W. Huang and B. von Stengel, “Computing an extensive-form correlated equilibrium in polynomial time,” in *International Workshop on Internet and Network Economics*, 2008, pp. 506–513.