Notes on Sparse-Support Approximate Nash Equilibria

Hu Fu

October 17, 2018

We present a result by Lipton et al. (2003) on the existence of sparsely supported approximate Nash equilibria.

We wlil make use of the following Chernoff-Hoeffding Inequality. In the list of exercises we gave a guided proof of it. You may also refer to the handouts on these bounds.

Theorem 1 (Chernoff-Hoeffding Inequality)**.** *Let $X_1, \cdots X_n$ be independently drawn random variables, each of which takes value in $[0, 1]$ with probability 1. Let X be their average, $\frac{1}{n}\sum_i X_i$. For any $\epsilon > 0$, we have*

$$\mathbf{Pr}\left[X - \mathbf{E}\left[X\right] \geq \epsilon\right] \leq e^{-2n\epsilon^2};$$
$$\mathbf{Pr}\left[X - \mathbf{E}\left[X\right] \leq \epsilon\right] \leq e^{-2n\epsilon^2}.$$

We focus on a two player game where the action sets of the two players, A_1 and A_2, are both of size m. Recall that we use $u_i : A_1 \times A_2 \to \mathbb{R}$ to denote the utility function of player i. Nash has shown that a Nash equilibrium is guaranteed to exist. The following theorem shows that, if we relax the solution concept and require only that the players cannot unilaterally improve their utilities *by too much*, then there is always an equilibrium with sparse support, i.e., the number of actions played with positive probability is logarithmic in m.

Definition 1. For $\epsilon > 0$, a strategy profile $(\boldsymbol{s}_1, \boldsymbol{s}_1) \in \Delta(A_1) \times \Delta(A_2)$ is an ϵ-approximate Nash equilibrium if, for either player i and any deviation $a \in A_i$, we have $u_i(\boldsymbol{s}_i, \boldsymbol{s}_{-i}) \geq u_i(a, \boldsymbol{s}_{-i}) - \epsilon$.

Without loss of generality, we will normalize the utilities so that $u_i(a_1, a_2) \in [0, 1]$, for any $i, a_1 \in A_1, a_2 \in A_2$.

Theorem 2 (Lipton et al., 2003)**.** *For any $\epsilon \in (0, 1)$, there exists an ϵ-approximate Nash equilibrium in which each player plays $O(\log m/\epsilon^2)$ strategies with positive probability.*

The proof is an instance of the probabilistic method, proving the existence of an object by showing that the probability of its occurrence is strictly positive. The idea is that working on a large enough set of samples i.i.d. drawn from a distribution is not much different from working on the original distribution — the very same idea as for Empirical Risk Minimization in machine learning. The bulk of the argument is to calculate the necessary size of the sample set to approximate the original distribution.

Proof. By Nash's theorem, a Nash equilibrium always exists. Let $(\boldsymbol{s}_1^*, \boldsymbol{s}_2^*)$ be a Nash equilibrium.

Take k i.i.d. samples, $a_{i1}, \ldots, a_{ik}$, from A_i, according to the distribution $\boldsymbol{s}_i^*$, for each $i \in \{1, 2\}$. Let $\tilde{\boldsymbol{s}}_i$ be the "empirical" strategy, which plays one of the k sampled actions uniformly at random.

We will show that, when k is $\Omega(\log m/\epsilon^2)$, with positive probability, for each player i and action $a \in A_i$, $u_i(a, \boldsymbol{s}^*_{-i}) - \frac{\epsilon}{2} \le u_i(a, \tilde{\boldsymbol{s}}_{-i}) \le u_i(a, \boldsymbol{s}^*_{-i}) + \frac{\epsilon}{2}$, which implies,

$$u_i(a, \tilde{\boldsymbol{s}}_{-i}) \le u_i(a, \boldsymbol{s}^*_{-i}) + \frac{\epsilon}{2} \le \frac{1}{k}\sum_{j=1}^{k} u_i(a_{ij}, \boldsymbol{s}^*_{-i}) + \frac{\epsilon}{2} \le \frac{1}{k}\sum_{j} u_i(a_{ij}, \tilde{\boldsymbol{s}}_{-i}) + \epsilon = u_i(\tilde{\boldsymbol{s}}_i, \tilde{\boldsymbol{s}}_{-i}) + \epsilon. \quad (1)$$

(In the second inequality we used the fact that any action played with positive probability in a Nash equilibrium gives no less utility than any other action against the opponent's Nash strategy.) In other words, with positive probability, $(\tilde{\boldsymbol{s}}_1, \tilde{\boldsymbol{s}}_2)$ constitutes an ϵ-approximate Nash; since $\tilde{\boldsymbol{s}}_i$ is supported on k actions, this proves the theorem.

In the sequel we focus on bounding the probability with which $u_i(a, \tilde{\boldsymbol{s}}_{-i})$ differs from $u_i(a, \boldsymbol{s}^*_{-i})$ by more than $\epsilon/2$, for any $a \in A_i$. Note that $u_i(a, \tilde{\boldsymbol{s}}_{-i}) = \frac{1}{k}\sum_{j=1}^{k} u_i(a, a_{-i,j})$; as $a_{-i,j}$ is a random action drawn from the distribution $\boldsymbol{s}^*_{-i}$, the expectation of $u_i(a, a_{-i,j})$ is just $u_i(a, \boldsymbol{s}^*_{-i})$. Therefore $u_i(a, \tilde{\boldsymbol{s}}_{-i})$ is the average of k i.i.d. random variables, each taking values from $[0, 1]$, with expectation equal to $u_i(a, \boldsymbol{s}^*_{-i})$. We are interested in the probability with which this average deviates from its expectation by a margin of $\frac{\epsilon}{2}$. This is precisely the situation to which the Hoeffding Inequality applies.

Apply Hoeffding inequality, we have

$$\mathbf{Pr}\left[|u_i(a, \tilde{\boldsymbol{s}}_{-i}) - u_i(a, \boldsymbol{s}^*_{-i})| > \frac{\epsilon}{2}\right] \le 2e^{-k\epsilon^2/2}.$$

Let's call the event $|u_i(a, \tilde{\boldsymbol{s}}_{-i}) - u_i(a, \boldsymbol{s}^*_{-i})| > \frac{\epsilon}{2}$ the bad event for action $a \in A_i$. There are $2n$ such events as a ranges over $A_1 \cup A_2$. Using the union bound, the probability that *any* of these $2n$ events happens is at most $4ne^{-k\epsilon^2/2}$. For $k > 2\log(4n)/\epsilon^2$, this is smaller than 1, and therefore with positive probability, *none* of the bad events happen. (1) shows that the corresponding $(\tilde{\boldsymbol{s}}_1, \tilde{\boldsymbol{s}}_2)$ is an ϵ-approximate Nash equilibrium. □

Remark 1. This existence theorem also gives rise to an algorithm for computing an approximate Nash equilibrium: by simply enumerating the sample sets and verifying whether the resulting "empirical" strategies constitute an approximate Nash equilibrium. This runs in time $m^{O(\log m/\epsilon^2)}$. This is called a *quasi-polynomial time* algorithm. There are various faster algorithms for special games. For example, Barman (2018) showed that, if in a two-player game with m actions for each player, any column of sum of their utility matrices contains no more than s nonzero entries, then there is an algorithm computing an ϵ-approximate Nash in time $m^{O(\log s/\epsilon^2)}$. Alon et al. (2013) gave a polynomial-time approximating scheme (PTAS) for the case when the sum of utility matrices has logarithmic rank. In general, however, the quasi-polynomial time algorithm of LLM is asymptotically best possible, if certain complexity hypothesis true: Rubinstein (2016), in his tour de force paper, showed that there is an $\epsilon > 0$ for which computing an ϵ-approximate Nash equilibrium for two-player games needs time $m^{\log^{1-o(1)} m}$, assuming the Exponential Time Hypothesis for PPAD.

Remark 2. We showed the theorem for two players, but it is straightforward to extend it to similar statements for general games with more players.

References

Alon, N., Lee, T., Shraibman, A., and Vempala, S. (2013). The approximate rank of a matrix and its algorithmic applications: approximate rank. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 675–684.

Barman, S. (2018). Approximating nash equilibria and dense subgraphs via an approximate version of carathéodory's theorem. *SIAM J. Comput.*, 47(3):960–981.

Lipton, R. J., Markakis, E., and Mehta, A. (2003). Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, pages 36–41. ACM.

Rubinstein, A. (2016). Settling the complexity of computing approximate two-player nash equilibria. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 258–265.