Approximating Nash Equilibria in Bimatrix Games

Yiannis Tselekounis

June 23, 2011

Notation

- For an integer n, $[n] = \{1, 2, \ldots, n\}$.
- $\mathbf{x}$: vector, x_i : the components of $\mathbf{x}$, $\mathbf{x}^T$: the transpose of $\mathbf{x}$.
- $\mathbb{P}^n$: the set of all probability vectors in n dimensions, i.e. $\mathbb{P}^n \equiv \{\mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^{n} x_i = 1,\ x_i \geq 0 \texttt{ for all } i \in [n]\}$.
- $\mathbb{R}^{n\times m}_{[0:1]}$: the set of all $n \times m$ matrices with real entries between 0 and 1, i.e. $\mathbb{R}^{n\times m}_{[0:1]} \equiv \{A \in \mathbb{R}^{n\times m}_{[0:1]} : 0 \leq \alpha_{i,j} \leq 1 \texttt{ for all } i \in [n],\ j \in [m]\}$.

Bimatrix Games

- 2-player games in which the set of strategies are S_1, S_2 respectively, $|S_1| = n$, $|S_2| = m$.
- Such games are denoted by $\Gamma = \langle A, B \rangle$ where A, $B \in \mathbb{R}^{n \times m}$.
- n rows of A, B represent the pure strategies of row player, m columns represent the pure strategies of the column player.
- For the strategies i (row) ,j (column), the row player gets payoff a_{ij} and the colum player gets b_{ij}.
- A mixed strategy for $i \in N$ is a probability distribution on the set of her pure strategies S_i ($\mathbf{x} \in \mathbb{P}^n$ for the row player,$\mathbf{y} \in \mathbb{P}^m$ for the column player).
- For mixed strategies $\mathbf{x}$ (row), $\mathbf{y}$ (column), the expected payoffs are $\mathbf{x}^T A\mathbf{y}$, $\mathbf{x}^T B\mathbf{y}$ respectively.

Nash equilibria

Definition 1 (Nash Equilibrium)

$(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ *is a Nash equilibrium for* $\Gamma = \langle A, B \rangle$ *if*

1. $\mathbf{x}^T A \tilde{\mathbf{y}} \leq \tilde{\mathbf{x}}^T A \tilde{\mathbf{y}}$, $\forall \mathbf{x} \in \mathbb{P}^n$ *and*
2. $\tilde{\mathbf{x}}^T B \mathbf{y} \leq \tilde{\mathbf{x}}^T B \tilde{\mathbf{y}}$, $\forall \mathbf{y} \in \mathbb{P}^m$.

Definition 2 (ϵ-Nash Equilibrium)

$(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ *is an* ϵ*-Nash equilibrium for* $\Gamma = \langle A, B \rangle$ *if*

1. $\mathbf{x}^T A \hat{\mathbf{y}} \leq \hat{\mathbf{x}}^T A \hat{\mathbf{y}} + \epsilon$, $\forall \mathbf{x} \in \mathbb{P}^n$ *and*
2. $\hat{\mathbf{x}}^T B \mathbf{y} \leq \hat{\mathbf{x}}^T B \hat{\mathbf{y}} + \epsilon$, $\forall \mathbf{y} \in \mathbb{P}^m$.

Definition 3 (ϵ-well supported Nash Equilibrium)

$(\mathbf{x}^*, \mathbf{y}^*)$ *is an* ϵ*-Nash equilibrium for* $\Gamma = \langle A, B \rangle$ *if*

1. $\forall i : x_i^* > 0 \Rightarrow e_j^T A \mathbf{y}^* \leq e_i^T A \mathbf{y}^* + \epsilon, \forall j$
2. $\forall i : y_i^* > 0 \Rightarrow \mathbf{x}^{*T} B e_j \leq \mathbf{x}^{*T} B e_i + \epsilon, \forall j$

Positively normalized bimatrix games

1. - Consider the game $\Gamma = \langle A, B \rangle$ and let c, d be two positive real constants.
   - Suppose that $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is a Nash equilibrium for Γ and $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is an ϵ-Nash equilibrium for Γ.
   - Consider the game $\Gamma' = \langle cA, dB \rangle$. Then Γ, Γ' have the same set of Nash Equilibria and any ϵ-Nash equilibrium for Γ is a $\lambda\epsilon$-Nash equilibrium for Γ' ($\lambda = \max\{c, d\}$).
2. Let C, $D \in \mathbb{R}^{m \times n}$ such that
   - for all (columns) $j \in [m]$, $c_{i,j} = c_j \in \mathbb{R}$ for all $i \in [n]$.
   - for all (rows) $i \in [m]$, $d_{i,j} = d_i \in \mathbb{R}$ for all $j \in [m]$.

   If we consider the game $\Gamma'' = \langle A + C, B + D \rangle$, then Γ, Γ'' are equivalent as regards their sets of Nash and ϵ-Nash equilibria.

1, 2 allows us to focus only on bimatrix games where the payoffs are between 0 and 1, i.e. on games A, B where A, $B \in \mathbb{R}^{m \times n}_{[0:1]}$ (*positively normalized bimatrix games*).

Related Work

Theorem 1

For any Nash equilibrium $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ of a positively normalized $n \times n$ bimatrix game and for every $\epsilon > 0$, there exists, for every $k \geq \frac{12 \ln n}{\epsilon^2}$, a pair of k-uniform strategies $\hat{\mathbf{x}}, \hat{\mathbf{y}}$ such that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is an ϵ-Nash equilibrium.

Theorem 2

The problem of computing a $\frac{1}{n^{\Theta(1)}}$-Nash equilibrium of a positive normalized $n \times n$ bimatrix game is PPAD-complete.

First Approximations

Lemma 1 (A $\frac{3}{4}$-Nash equilibrium)

Consider any positively normalized $n \times m$ bimatrix game $\Gamma = \langle A, B \rangle$ and let $\alpha_{i_1,j_1} = \max_{i,j}\ \alpha_{i,j}$ and $b_{i_2,j_2} = \max_{i,j}\ b_{i,j}$. Then the pair of strategies $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ where $\hat{x}_{i_1} = \hat{x}_{i_2} = \hat{y}_{j_1} = \hat{y}_{j_2} = \frac{1}{2}$ is a $\frac{3}{4}$-Nash equilibrium for Γ.

Theorem 3 (A Parameterized Approximation)

Consider a positively normalized $n \times m$ bimatrix game $\Gamma = \langle A, B \rangle$. Let λ_1^* (λ_2^*) be the minimum, among all Nash equilibria of Γ, expected payoff for the row (column) player and let $\lambda = \min\{\lambda_1^*, \lambda_2^*\}$. Then, for any $0 < \epsilon < 1$, there exists a $\frac{2+\lambda+\epsilon}{4}$ -Nash equilibrium that can be computed in time polynomial in $\frac{1}{\epsilon}$, n and m.

A Simple Algorithm for $\frac{1}{2}$-Nash equilibrium

1. Pick an arbitrary row for the row player, say row i.
2. Let $j = arg \max_j C_{ij}$.
3. Let $k = arg \max_k R_{kj}$.
4. The equilibrium is $\hat{\mathbf{x}} = \frac{1}{2}e_i + \frac{1}{2}e_k$, $\hat{\mathbf{y}} = e_j$.

Theorem 4

The strategy pair $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ is a $\frac{1}{2}$-Nash equilibrium.

Decorrelated Games

Definition 4

A mapping $f : [n] \to [n]$ *is a best response mapping for the column player iff, for every* $i \in [n]$, $C_{if(i)} = \max_j C_{ij}$.

Definition 5 (Decorrelation Transformation)

The decorrelated game (R^f, C^f) *corresponding to the best response mapping* f *is defined as follows* $\forall i, j \in [n]$:
$R^f_{ij} = R_{if(j)}, \quad C^f_{ij} = C_{if(j)}$.

Decorrelated Games (Cont.)

Lemma 2

In the game (R^f, C^f), for all sets $S \subseteq [n]$, the strategies of the column player in S are $\frac{|S|-1}{|S|}$-well supported against the strategy $\mathbf{x}^*$ of the row player, where $\mathbf{x}^*$ is defined in terms of the set $S' = \{i \in S | C^f_{ii} = 0\}$ as follows

- *if $S' \neq \emptyset$, then $\mathbf{x}^*$ is uniform over the set S'*
- *if $S' = \emptyset$, then*

$$\mathbf{x}^*_i = \begin{cases} \frac{1}{Z}\frac{1}{C^f_{ii}}, & \textit{if } i \in S \\ 0, & \textit{otherwise} \end{cases}$$

where $Z = \sum_{i \in S} \frac{1}{C^f_{ii}}$ is a normalizing constant.

Decorrelated Games (Cont.)

Lemma 3 (Player Decorrelation)

In the game (R^f, C^f)*, if there exists a set* $S \subseteq [n]$ *and a mixed strategy* $\mathbf{y} \in \Delta(S)$ *for the column player such that the strategies in* S *are* $\frac{|S|-1}{|S|}$*-well supported for the row player against the distribution* $\mathbf{y}$*, then there exists a strategy* $\mathbf{x} \in \Delta(S)$ *for the row player so that the pair* $(\mathbf{x}, \mathbf{y})$ *is an* $\frac{|S|-1}{|S|}$*-well supported Nash equilibrium.*

Decorrelated Games (Cont.)

Lemma 4

For all $S \subseteq [n]$, if the pair $(\mathbf{x}^*, \mathbf{y}^*)$, where $\mathbf{x}^*$ is defined as in the statement of Lemma 2 and $\mathbf{y}^*$ is the uniform distribution over S, constitutes an $\frac{|S|-1}{|S|}$-well supported Nash equilibrium for the game (R^f, C^f), then the pair of distributions $(\mathbf{x}^*, \mathbf{y}')$ is an $\frac{|S|-1}{|S|}$-well supported Nash equilibrium for the game (R, C), where y' is the distribution defined as follows

$$y'(i) = \sum_{j \in S} \mathbf{y}^*(j) X_{f(j)=i}, \forall i \in [n],$$

where $X_{f(j)=i}$ is the indicator function of the condition "$f(j) = i$".

Reduction to Win-Lose Game

We define a mapping from a general 2-player game to a win-lose game so that well supported equilibria of the win-lose game can be mapped to well supported equilibria of the original game.

$$round(A)_{ij} = \begin{cases} 1, & \text{if } A_{ij} \geq \frac{1}{2} \\ 0, & \text{if } A_{ij} < \frac{1}{2} \end{cases}$$

Lemma 5

If $(\mathbf{x}, \mathbf{y})$ is an ϵ-well supported Nash equilibrium of the game $(round(R), round(C))$, then $(\mathbf{x}, \mathbf{y})$ is a $\frac{1+\epsilon}{2}$-well supported Nash equilibrium of the game (R, C).

Algorithm for well supported equilibria

ALG-WS

1. Map game (R, C) to the win-lose game $(round(R), round(C))$.
2. Map game $(round(R), round(C))$ to the game $(round(R)^f, round(C)^f)$, where f is any best response mapping for the column player.
3. Find a subset $S \subseteq [n]$ and a strategy $\mathbf{y} \in \Delta(S)$ for the column player such that all the strategies in S are $\frac{|S|-1}{|S|}$-well supported for the row player in $(round(R)^f, round(C)^f)$ against the strategy $\mathbf{y}$ for the column player.
4. By a successive application of lemmas 3, 4 and 5, get an $\frac{1}{2} + \frac{1}{2}\frac{|S|-1}{|S|} = 1 - \frac{1}{2|S|}$ well supported Nash equilibrium of the original game.

A non-trivial step

Step 3

"Given a $0-1$ matrix $round(R)^f$, find a subset of the columns $S \subseteq [n]$ and a distribution $\mathbf{y} \in \Delta(S)$, so that all rows in S are $\frac{|S|-1}{|S|}$-well supported against the distribution $\mathbf{y}$ over the columns."

Conjecture 1

There are integers κ and λ such that every digraph either has a cycle of length at most κ or an undominated set of λ vertices.

Theorem 5

If Conjecture 1 is true for some values of κ and λ, then Algorithm ALG-WS returns in polynomial time (e.g. by exhaustive search) a $\max\{1-\frac{1}{2\kappa}, 1-\frac{1}{2\lambda}\}$-well-supported Nash equilibrium which has support of size $\max\{\kappa, \lambda\}$.

A simpler approach

Theorem 6

For arbitrary win lose bimatrix game A, B , there is a polynomial time constructible profile that is a 0.5*-well supported Nash equilibrium of the game.*

Corollary 1

For any $[0,1]$*-bimatrix game R, C , there is a* 0.75*-well supported Nash equilibrium that can be computed in polynomial time.*

Improved Approximations

[DMP]

An approximation ratio of $0.38 + \epsilon$, $\forall \epsilon > 0$, based on the following

1. If the values (u, v) of a Nash equilibrium to the two players were known, then we would be able to find a $\max\{u, v\}$ -approximate Nash equilibrium by solving a set of linear inequalities.
2. For every Nash equilibrium, there is a pair of mixed strategies with support size $O(\frac{1}{\epsilon^2})$ which approximates within ϵ the true values of that Nash equilibrium.

[BBM]

An approximation ratio of 0.36 based on solving a zero-sum game.

Thank you!