The Multilinear Minimax Relaxation of Bimatrix Games and Comparison with Nash Equilibria via Lemke-Howson

Chun Leung Lau (Jointly with B. Kalantari)

Department of Computer Science, Rutgers University, New Brunswick, New Jersey, USA

30th International Conference on Economics
July 16, 2019

Outline

- Von Neumann Minimax Theorem (Two Players)

Outline

- Von Neumann Minimax Theorem (Two Players)
- Nash Equilibria in the Bimatrix Case

Outline

- Von Neumann Minimax Theorem (Two Players)
- Nash Equilibria in the Bimatrix Case
- Explanation of MMR as an Approximation of Nash Equilibria

Outline

- Von Neumann Minimax Theorem (Two Players)
- Nash Equilibria in the Bimatrix Case
- Explanation of MMR as an Approximation of Nash Equilibria
- Experimental Results

Outline

- Von Neumann Minimax Theorem (Two Players)
- Nash Equilibria in the Bimatrix Case
- Explanation of MMR as an Approximation of Nash Equilibria
- Experimental Results
- Balanced Scaling

Unit Simplex

For a positive integer k denote the unit simplex by

$$S_k = \{u \in \mathbb{R}^k : \sum_{j=1}^{k} u_j = 1, u \geq 0\}. \quad (1)$$

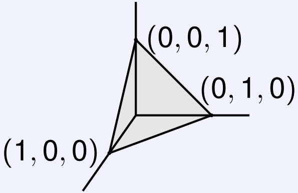

Figure: The unit simplex S_k.

Von Neumann Minimax Theorem - Zero-Sum Games

Theorem

Let $A = (a_{ij})$ be an $m \times n$ real matrix.

$$\max_{x \in S_m} \min_{y \in S_n} x^T A y = \min_{y \in S_n} \max_{x \in S_m} x^T A y,$$

where S_k is the unit simplex in $\mathbb{R}^k$, i.e.

$$S_k = \{x \in \mathbb{R}^k : \quad \sum_{i=1}^{k} x_k = 1, \quad x \geq 0\}.$$

Nash Equilibria - Non-Cooperative Games

Theorem

(Nash) *Let* $R = (R_{ij})$ *and* $C = (C_{ij})$ *be* $m \times n$ *real payoff matrices. Given* $p = (x, y) \in S_m \times S_n$*, set*

$$R[p] = x^T Ry, \quad C[p] = x^T Cy.$$

There exists a composite vector $p_* = (x_*, y_*)$ *(called Nash Equilibrium) such that*

$$\max\{R[p] : p = (x, y_*),\ x \in S_m\} = R[p_*]$$

and

$$\max\{C[p] : p = (x_*, y),\ y \in S_n\} = C[p_*].$$

Nash Equilibria - Non-Cooperative Games

Theorem

(Nash) *Let* $R = (R_{ij})$ *and* $C = (C_{ij})$ *be* $m \times n$ *real payoff matrices. Given* $p = (x, y) \in S_m \times S_n$, *set*

$$R[p] = x^T R y, \quad C[p] = x^T C y.$$

There exists a composite vector $p_* = (x_*, y_*)$ *(called Nash Equilibrium) such that*

$$\max\{R[p] : p = (x, y_*),\ x \in S_m\} = R[p_*]$$

and

$$\max\{C[p] : p = (x_*, y),\ y \in S_n\} = C[p_*].$$

It is not known if p_* can be computed efficiently.

The Multilinear Minimax Relaxation (MMR) Theorem

Theorem

(Kalantari) Set $\Delta = S_m \times S_n$*. Given* $\alpha = (\alpha_1, \alpha_2) \in S_2$ *and* $p = (x, y) \in \Delta$*, let*

$$A_0[\alpha, p] = \alpha_1 R[p] + \alpha_2 C[p], \tag{2}$$

Then,

$$\min_{\alpha \in S_2} \max_{p \in \Delta} A_0[\alpha, p] = \max_{p \in \Delta} \min_{\alpha \in S_2} A_0[\alpha, p]. \tag{3}$$

Let α^* *and* $p^* = (x^*, y^*)$ *be a minimax solution. Then, for any* $(\alpha, p) \in S_2 \times \Delta$ *we have*

$$A_0[\alpha^*, p] \leq A_0[\alpha^*, p^*] \leq A_0[\alpha, p^*] \tag{4}$$

Furthermore, (α^*, p^*) *is computable via a primal-dual pair of linear programming with* $O(m \times n)$ *variables.* □

A Game-Theoretic Interpretation of MMR

The game is played as a meta-player against *Row* and *Column*.

The meta-player chooses an $\alpha \in S_2$. *Row* chooses $x \in S_m$ and *Column* chooses $y \in S_n$.

Once α and $p = (x, y)$ are chosen, the meta-player's expected loss is $A_0[\alpha, p]$, while *Row* and *Column* gains the same amount.

Under pure strategies, the meta-player plays against either *Row* and *Column*.

Applications to Approximations of Nash Equilibria

Corollary

Let p_* be a Nash equilibrium and p^* a solution to MMR. Then

$$\alpha_1^* R[p_*] + \alpha_2^* C[p_*] \leq \alpha_1^* R[p^*] + \alpha_2^* C[p^*].$$

In other words we can approximate a convex combination of R and C payoffs under Nash equilibrium.

Corollary

Either $R[p_*] \leq R[p^*]$, or $C[p_*] \leq C[p^*]$. In particular, if both inequalities are satisfied, then p^* gives a better payoff than p_* for both players. □

Algorithmic Computation of MMR

The primal-dual LP formulation of MMR is the following:

$$
\begin{aligned}
& \min \delta \\
(LP) \qquad & R_{ij}\alpha_1 + C_{ij}\alpha_2 \leq \delta, \quad \forall i \in [1, n], j \in [1, m], \\
& \alpha_1 + \alpha_2 = 1, \\
& \alpha_1 \geq 0, \\
& \alpha_2 \geq 0.
\end{aligned}
$$

Algorithmic Computation of MMR

$$
(DLP) \qquad
\begin{aligned}
&\max \lambda \\
&\sum_{i=1}^{n}\sum_{j=1}^{m} R_{ij} q_{ij} \geq \lambda, \\
&\sum_{i=1}^{n}\sum_{j=1}^{m} C_{ij} q_{ij} \geq \lambda, \\
&\sum_{i=1}^{n}\sum_{j=1}^{m} q_{ij} = 1, \\
&q_{ij} \geq 0, \quad \forall i \in [1, n], j \in [1, m].
\end{aligned}
$$

Algorithmic Computation of MMR

Let Q be the $m \times n$ matrix (q_{ij}). Given a positive integer k let $e_k = (1, \ldots, 1)^T \in \mathbb{R}^k$. Then set

$$x^* = Qe_n \quad y^* = Q^T e_m.$$

Then x^* and y^* are the strategies for *Row* and *Column*.

Computational Results

Using Gambit 15.1.1, the entries of payoff matrices are generated randomly between 0 and 1.

Computational Results

Using Gambit 15.1.1, the entries of payoff matrices are generated randomly between 0 and 1.

We compared Lemke-Howson algorithm with MMR algorithm in the LP formulation.

Computational Results

Using Gambit 15.1.1, the entries of payoff matrices are generated randomly between 0 and 1.

We compared Lemke-Howson algorithm with MMR algorithm in the LP formulation.

Programs run for a maximum time of 30 minutes before they are terminated and returned an error. For each comparison, we ran the algorithms 1000 times.

Computational Results

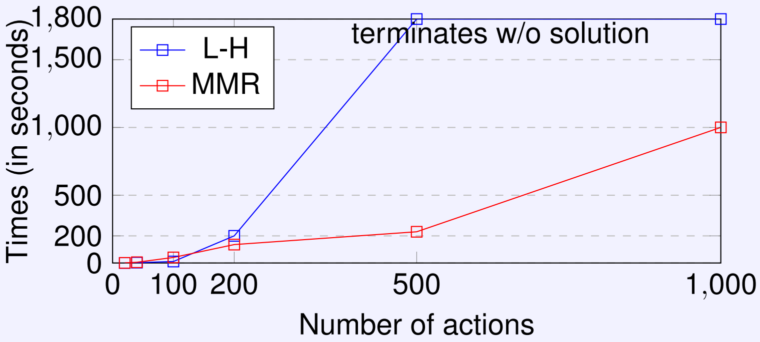

Figure: Speed comparisons for Lemke-Howson and MMR algorithms

Computational Results

Actions	Percentages
10	81.2%
30	80.3%
60	79.98%
100	82.0055%
150	80.023%

Figure: Percentages of times MMR algorithm had better payoffs for both players

Computational Results

Actions	Relative errors
10	$0.129 \pm .020$
30	$0.090 \pm .012$
60	$0.110 \pm .010$
100	$0.151 \pm .016$
150	$0.141 \pm .004$

Figure: Relative errors of payoffs if one player's MMR payoff did worse than the Nash equilibrium payoff

Computational Results

Actions	Relative errors
10	$0.129 \pm .020$
30	$0.090 \pm .012$
60	$0.110 \pm .010$
100	$0.151 \pm .016$
150	$0.141 \pm .004$

Figure: Relative errors of payoffs if one player's MMR payoff did worse than the Nash equilibrium payoff

As shown by the experimental results, MMR algorithm outperforms Lemke-Howson both in speed and the quality of payoffs.

Balanced Scaling Theorem

Definition

A profile is called *balanced* if both expected payoffs are equal.

Balanced Scaling Theorem

Definition

A profile is called *balanced* if both expected payoffs are equal.

Theorem

Given payoff matrices R, C $\in [0,1]^{n\times n}$*, for each t* $\in [0,1]$ *let* $p^*(t) = (p^{1*}(t), p^{2*}(t))$ *be the MMR solution corresponding to the case when R is replaced by tR and C by* $(1-t)C$*. Then there exists* $t_* \in [0,1]$ *that gives a balanced scaling, i.e.*

$$t_* R[p^*(t_*)] = (1-t_*)C[p^*(t_*)], \tag{5}$$

Balanced Scaling Theorem

Definition

A profile is called *balanced* if both expected payoffs are equal.

Theorem

Given payoff matrices R, C $\in [0,1]^{n\times n}$, *for each* $t \in [0,1]$ *let* $p^*(t) = (p^{1*}(t), p^{2*}(t))$ *be the MMR solution corresponding to the case when R is replaced by tR and C by* $(1-t)C$. *Then there exists* $t_* \in [0,1]$ *that gives a balanced scaling, i.e.*

$$t_* R[p^*(t_*)] = (1-t_*)C[p^*(t_*)], \tag{5}$$

To compute a convex scaling that gives balanced MMR payoffs, we can use the Bisection Method to change the value of t.

Balanced Scaling

The advantage is that there exists a *convex* scaling pair $d = (t, 1 - t) \in S_2$ so that the MMR expected payoffs for *Row* and *Column* players are equal.

At least one player's expected payoff is at least as good as the expected payoffs of any Nash equilibria.

Visualizations of Balanced Scalings

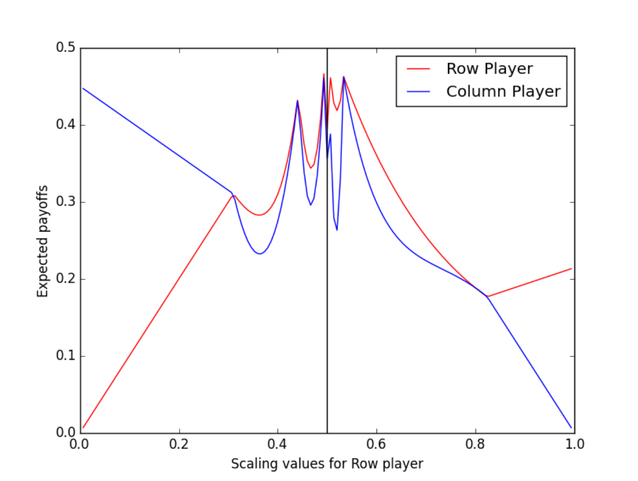

Summary

- Based on extensive computational results, MMR is an effective approximation to Nash Equilibria.

Summary

• Based on extensive computational results, MMR is an effective approximation to Nash Equilibria.

• In a good percentage of time, it out-performs Lemke-Howson in speed and in quality of solutions.

Summary

• Based on extensive computational results, MMR is an effective approximation to Nash Equilibria.

• In a good percentage of time, it out-performs Lemke-Howson in speed and in quality of solutions.

• Because Nash equilibria can be applied to machine-learning algorithms and robotics, we expect the applicability of MMR in such problems.

Related Articles

B. Kalantari, Approximating Nash Equilibrium Via Multilinear Minimax, arxiv.org/pdf/1605.00167.pdf, 2019.

B. Kalantari and C. L. Lau, The Multilinear Minimax Relaxation of Bimatrix Games and Comparison with Nash Equilibria Via Lemke-Howson, arxiv.org/pdf/1809.01717.pdf, 2019.

THANK YOU!