Bayesian Action-Graph Games

Albert Xin Jiang and Kevin Leyton-Brown
Department of Computer Science
University of British Columbia

November 13, 2011

Equilibrium Computation in Bayesian Games

Equilibrium Computation in Bayesian Games

Games of incomplete information (Bayesian games)

- Proposed by Harsanyi (1967)
- Players are uncertain about game being played
- Each player receive private information (type)
- Many applications in economics: e.g. auctions

Equilibrium Computation in Bayesian Games

Games of incomplete information (Bayesian games)

- Proposed by Harsanyi (1967)
- Players are uncertain about game being played
- Each player receive private information (type)
- Many applications in economics: e.g. auctions

Bayesian Game $(N, \{A_i\}_{i \in N}, \Theta, P, \{u_i\}_{i \in N})$

- set of players $N = \{1, \ldots, n\}$
- each player i's action set: A_i
- set of type profiles $\Theta = \prod_i \Theta_i$
- type distribution $P : \Theta \to \mathbb{R}$

Equilibrium Computation in Bayesian Games

Games of incomplete information (Bayesian games)

- Proposed by Harsanyi (1967)
- Players are uncertain about game being played
- Each player receive private information (type)
- Many applications in economics: e.g. auctions

Bayesian Game $(N, \{A_i\}_{i \in N}, \Theta, P, \{u_i\}_{i \in N})$

- set of players $N = \{1, \dots, n\}$
- each player i's action set: A_i
- set of type profiles $\Theta = \prod_i \Theta_i$
- type distribution $P : \Theta \to \mathbb{R}$
- player i's utility function $u_i : A \times \Theta \to \mathbb{R}$

Bayes-Nash Equilibrium

- Natural extension of Nash equilibrium to Bayesian games

Bayes-Nash Equilibrium

- Natural extension of Nash equilibrium to Bayesian games
- Mixed strategy θ_i
  - probability of playing action a_i given type θ_i is $\theta_i(a_i|\theta_i)$
- Expected utility of i given θ_i is
  $u_i(\sigma|\theta_i) = \sum_{\theta_{-i}} P(\theta_{-i}|\theta_i) \sum_a u_i(a, \theta) \prod_j \sigma_j(a_j|\theta_j)$

Bayes-Nash Equilibrium

- Natural extension of Nash equilibrium to Bayesian games
- Mixed strategy θ_i
  - probability of playing action a_i given type θ_i is $\theta_i(a_i|\theta_i)$
- Expected utility of i given θ_i is $u_i(\sigma|\theta_i) = \sum_{\theta_{-i}} P(\theta_{-i}|\theta_i) \sum_a u_i(a,\theta) \prod_j \sigma_j(a_j|\theta_j)$
- Mixed strategy profile σ is Bayes-Nash equilibrium if for all i, for all θ_i, for all a_i,

$$u_i(\sigma|\theta_i) \geq u_i(\sigma^{\theta_i \rightarrow a_i}|\theta_i)$$

Obstacles to efficient computation

- Representation
  - The straightforward Bayesian Normal Form requires exponential space in number of players

Obstacles to efficient computation

- Representation
  - The straightforward Bayesian Normal Form requires exponential space in number of players
- Lack of practical algorithms
  - Can be reduced to finding a Nash equilibrium in a complete-information game
  - But this transformation causes a further exponential blowup in size

Compact Representations

Most games of interest have highly-structured utility functions

Compact Representations

Most games of interest have highly-structured utility functions

- Compact representations for complete-information games
  - Graphical games Kearns et al. 2001
  - Action-graph games Jiang et al. 2010

- Dynamic games
  - Multi-agent influence diagrams Koller & Milch 2001
  - Temporal action-graph games Jiang et al. 2009

Our Contributions

Bayesian Action-Graph Games (BAGGs)

- Can represent arbitrary Bayesian games
- Compactly express games with structure
  - symmetry/anonymity
  - action- and type- specific utility independence
  - probabilistic independence of type distribution

Our Contributions

Bayesian Action-Graph Games (BAGGs)

- Can represent arbitrary Bayesian games
- Compactly express games with structure
  - symmetry/anonymity
  - action- and type- specific utility independence
  - probabilistic independence of type distribution

- Efficient computation of Bayes-Nash equilibria
  - adapt existing algorithms for Nash equilibria
  - exponential speedup
  - software available http://agg.cs.ubc.ca

Bayesian Action-Graph Games

Represent type distribution P as a Bayesian network

- Containing at least n random variables representing $\theta_1, \ldots, \theta_n$

Bayesian Action-Graph Games

Represent type distribution P as a Bayesian network

- Containing at least n random variables representing $\theta_1, \ldots, \theta_n$

Represent utility functions on an action graph:

- directed graph on set of action nodes $\mathcal{A}$
- player i, given θ_i, chooses an action from type-action set $A_{i,\theta_i} \subseteq \mathcal{A}$
- for each action node α, action count: number of players that have chosen α

Bayesian Action-Graph Games

Represent type distribution P as a Bayesian network

- Containing at least n random variables representing $\theta_1, \ldots, \theta_n$

Represent utility functions on an action graph:

- directed graph on set of action nodes $\mathcal{A}$
- player i, given θ_i, chooses an action from type-action set $A_{i,\theta_i} \subseteq \mathcal{A}$
- for each action node α, action count: number of players that have chosen α
- utility depends only on action node chosen and the action counts of its neighbors

Simple Example

Symmetric Bayesian game, n players, 2 types, 2 actions per type

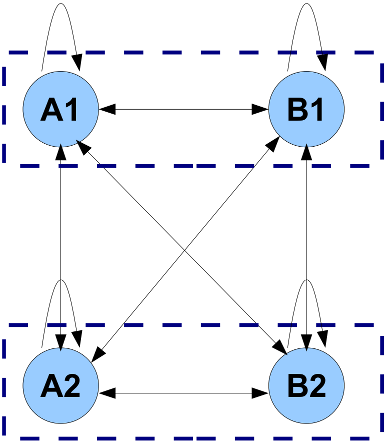

Type 1

Type 2

Theorem

if constant in-degrees, representation size is polynomial in n, $|\mathcal{A}|$, $|\Theta_i|$

Theorem

if constant in-degrees, representation size is polynomial in n, $|\mathcal{A}|$, $|\Theta_i|$

Extension: function nodes

- represents some function of its neighbors' action counts
- e.g. counting function node: sum

Coffee Shops

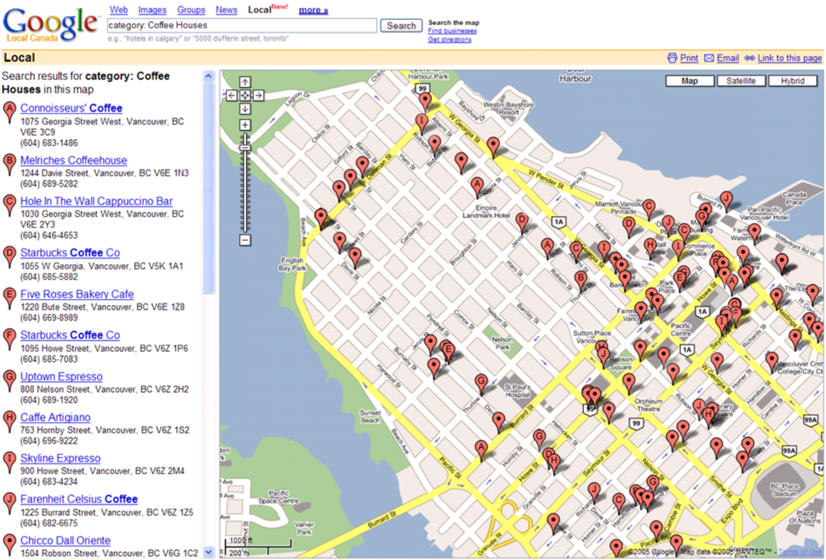

Example: Coffee Shop Game

- Each player chooses a location (in an r by k grid) to open a coffee shop, or decide not to enter.

- Utility of player i choosing a location depends on:
  - her type,
  - # of players choosing same block
  - # of players choosing surrounding blocks
  - # of players choosing any other block

Coffee Shop BAGG

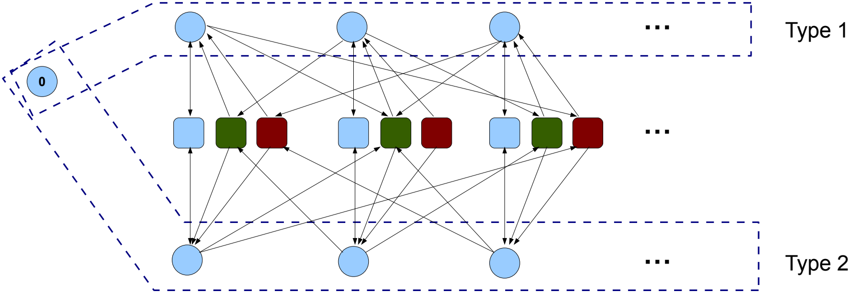

Computing Bayes-Nash Equilibria

Computing Bayes-Nash Equilibria

Reduce to complete-information game (agent form)

- one player for each type
- set of actions for player (i, θ_i): type-action set A_{i,θ_i}
- Nash equilibria correspond to Bayes-Nash of BAGG

Computing Bayes-Nash Equilibria

Reduce to complete-information game (agent form)

- one player for each type
- set of actions for player (i, θ_i): type-action set A_{i,θ_i}
- Nash equilibria correspond to Bayes-Nash of BAGG
- do not need to represent explicitly: the BAGG serves as a compact representation

Computing Bayes-Nash Equilibria (cont'd)

Adapt state-of-the-art algorithms for Nash equilibrium

- Global Newton Method Govindan & Wilson 2001
- Simplicial Subdivision van der Laan et al. 1987

A key subtask: computing expected utility (EU) of agent form given a mixed strategy profile

- equiv. to computing EU of the BAGG

Computing Bayes-Nash Equilibria (cont'd)

Adapt state-of-the-art algorithms for Nash equilibrium

- Global Newton Method Govindan & Wilson 2001
- Simplicial Subdivision van der Laan et al. 1987

A key subtask: computing expected utility (EU) of agent form given a mixed strategy profile

- equiv. to computing EU of the BAGG
- formulate as Bayesian network (BN) inference problem
- further exploit causal independence by creating intermediate variables

Computing EU in BAGGs

Theorem

for independent type distributions, EU can be computed in time polynomial in the size of the BAGG

Experiments: GW, Coffee Shop, 2 types, 6 locations

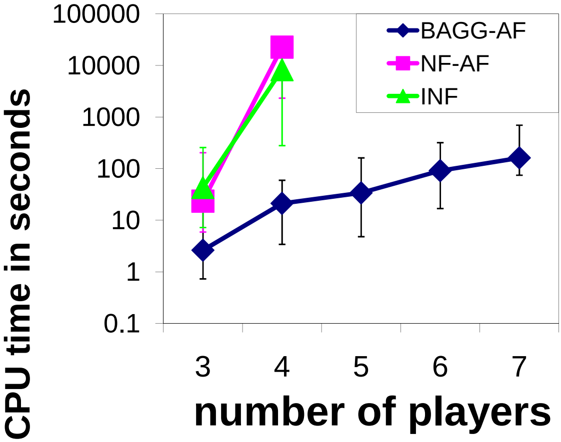

GW, Coffee Shop, 3 players, 3 locations

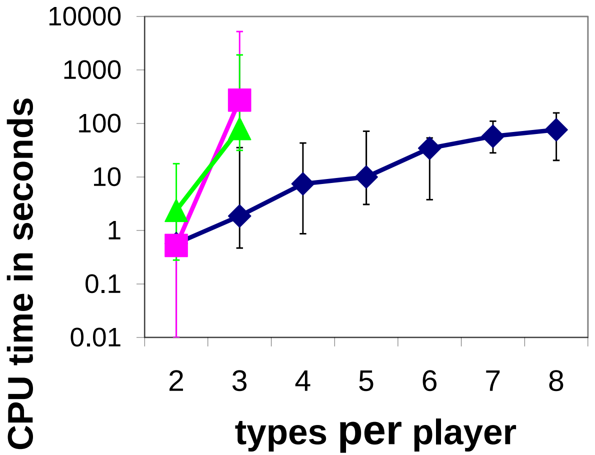

GW, Coffee Shop, 3 players, 2 types

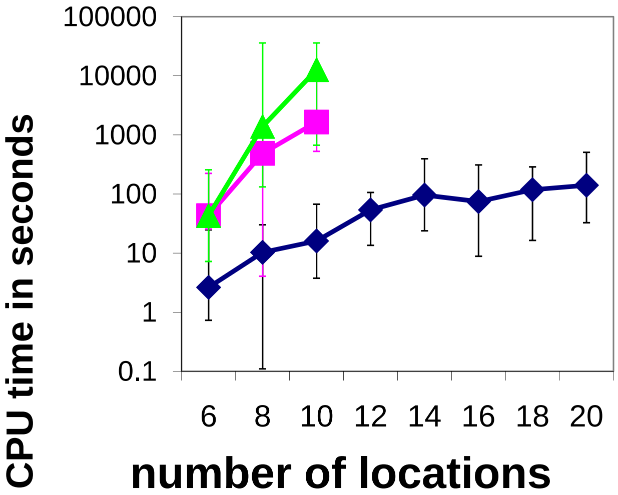

Simplicial Subdivision, Coffee Shop, 2 types, 3 locations

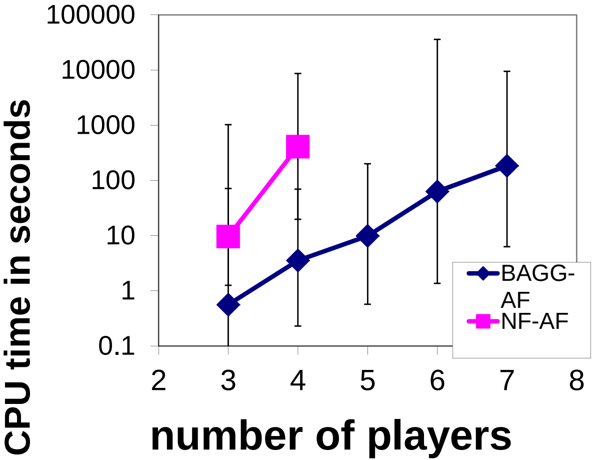

Conclusion

Bayesian Action-Graph Games

- exploit anonymity and action- and type-specific independence

Conclusion

Bayesian Action-Graph Games

- exploit anonymity and action- and type-specific independence

Computation

- compute Bayes-Nash equilibria by finding Nash equilibria in a complete-information game (agent-form)
- software available http://agg.cs.ubc.ca

Jiang, A.X., Leyton-Brown, K. Bayesian Action-Graph Games. In NIPS, 2010.