

A Language for Modeling Agents' Decision-making Processes in Games



Ya'akov Gal and Avi Pfeffer
Harvard University

Multi agent systems operate in the real world

- Players do not behave rationally.
- Players might be mistaken about game and opponents' beliefs.
- Players use reasoning processes outside of the game (heuristics, divide and conquer)
- Here, we need to distinguish between real-world and mental models used to deliberate.

Game theory and multi-agent systems

- Game theory has had a deep impact on AI.
- Game theory provides a language for decision-making under uncertainty.
- Given game description, agents' optimal decisions can be derived rationally (Nash equilibrium)

The traditional approach

- Bayesian Games (Harsanyi '68) can model uncertainty over decision-making processes.
- Each deliberation model is reduced to a discrete *type*.
- For each type, need to specify a separate
 - Game description.
 - Conditional probability distribution over opponents' types.
- Compute a Bayesian Nash equilibrium.

The problem with Bayesian Games

- They are a contrived representation for uncertainty over decision-making models.
 - Number of types for each agent is exponential in the number of decisions (e.g., player has n decisions and 2 models for each, need to define 2^{**n} types).
 - Must represent a joint distribution over all type instantiations.
 - Ill-suited for modeling heuristics.

Networks of Influence Diagrams (NIDs)

- Graphical language based on Multi Agent Influence Diagram (MAID) (Koller and Milch, 2001)
- MAIDs
 - Graphically describe a decision process where all agents know a correct model of the world.
 - Compute a Nash equilibrium (correct beliefs).

Contributions

- A *natural* and *compact* language for describing uncertainty over decision-making models.
- Models agents that are boundedly rational. (e.g., heuristics, divide-and-conquer)
- Compact representation –may be exponentially smaller than a Bayesian game.

Networks of Influence Diagrams (NIDs)

- Graphical language where each node is a Multi Agent Influence Diagram (MAID) (Koller and Milch, 2001)
- NIDs
 - Allow uncertainty over decision-making processes.
 - Compute an equilibrium given agents' beliefs, which might be incorrect !



The voting game


- Agents Black, Green and Red are voting one of them to chair a committee.
- Black will be chairman if voting draws.
- Agents Green and Red dislike Black and don't want him to chair.
- All agents prefer to win, but if they lose, they prefer to vote for the winner.

The voting game

- Possible pure Nash equilibria are for all agents to vote for the same person.
- Suppose Black wishes to model his opponents
 - If Green and Red vote the same, Black should vote for that agent.
 - If Green and Red miscoordinate, Black should vote for himself.
- Black is uncertain regarding the voting strategy of Green and Red.
- We can model this naturally using a NID.


Voting game NID

- A NID is a DAG where each node (or block) is a possible mental model.
- The leaves describe models in which both Red and Green are automata.
- The root block (top-level) represents the real-world from the modeler's point of view.



Quantifying belief over models


- At top-level block, Black believes :
 - Red and Green vote for Green (Red) with probability 0.6 (0.3).
 - Red and Green use top-level block with prob. 0.1



Solving NIDs bottom-up

At each block, output a best-response strategy for each agent given her beliefs

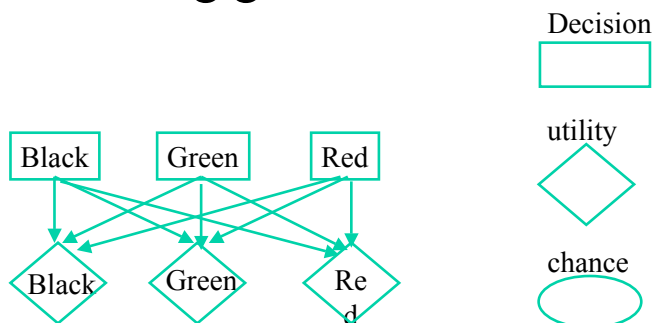
- Compute Nash equil. at the leaves.
- Eliminate node by passing strategy of modeled decision up to parent.
- Parent incorporates children's strategies probabilistically.



Describing relationship between models


- Suppose that Black believes that Red and Green might have colluded to vote for one of them.
- With some probability, Black believes Red or Green might renege.
- How do we capture this in a NID ?

Voting game MAID



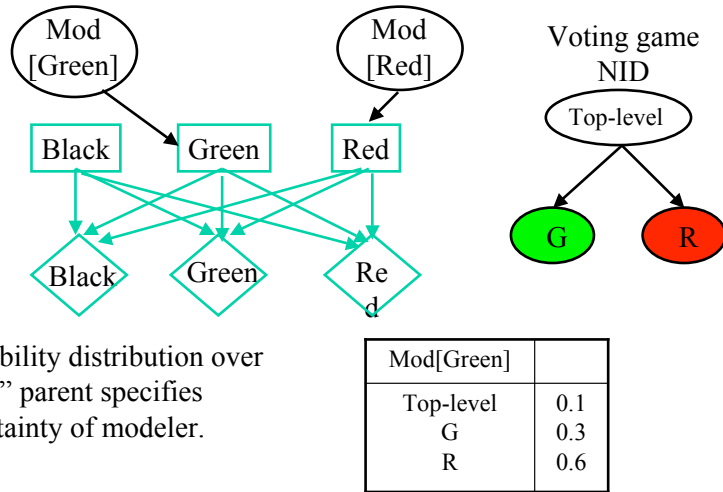
- Graph topology describes structure of game.
- Each node contains a table of parameters.

Quantifying beliefs in NIDs



- At each block, each decision has a “Mod” parent defining the block used to make the decision.
- When Mod(Green)= G, Black believes agent Green uses block G to make her decision.

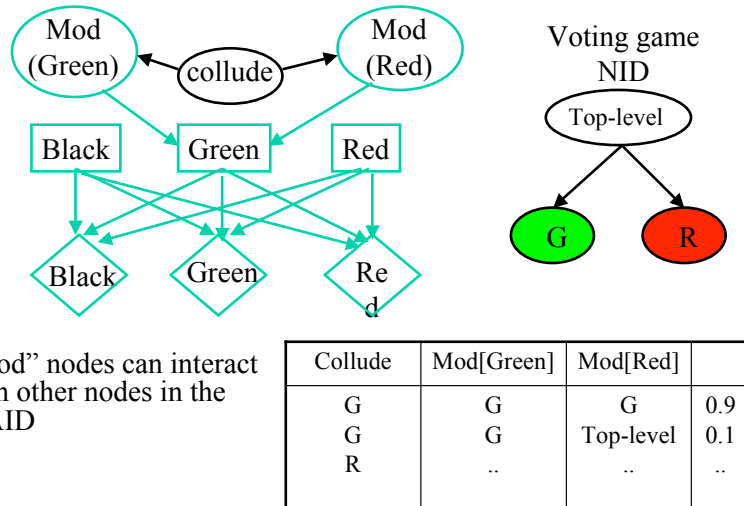
Quantifying beliefs in NIDs



Correlating between models

- Suppose that Black believes that Red and Green might collude to vote for one of them.
- With some probability, Red or Green might renege and play a Nash equilibrium.
- This can also be naturally captured in a NID.

The top-level MAID




Querying the NID

- We now have the ability to ask questions
 - What is Black's best-response strategy ?
 - What is the expected utility of each agent ?
 - What is the value of information of knowing whether Red and Green collude ?
 - Suppose Black's model is incorrect. What is the cost of the incorrect model ?

NIDs are compact

Player has n decisions, and modeler believes that each decision has 2 possible models.



Applications

- NIDs can be used for many purposes.
 - Modeling bounded rationality.
 - Opponent modeling.
 - See paper for details.

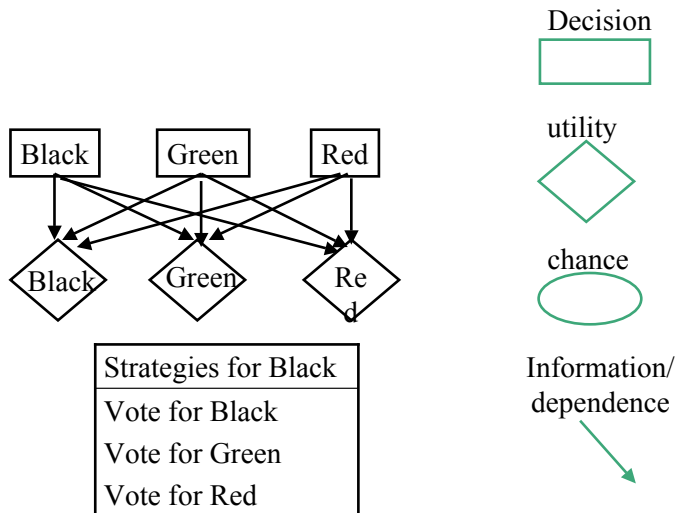
Current and future work

- Theoretical foundations.
 - Defining new type of equilibrium.
 - Solving cyclic NIDs.
 - Relationship with Bayesian games.
- Empirical investigations.
 - Using NIDs to model human-computer negotiation.
 - Learning opponents' models from data (Gal & Pfeffer, AAMAS GTDT 03)

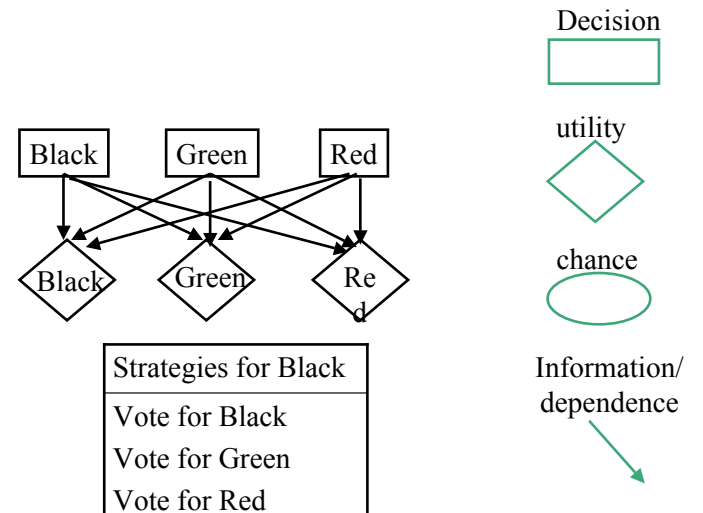
Conclusions

- To model uncertainty, must separate real-world from deliberation models.
- Traditional game theory is unsuitable for this.
- “Mod” variables in NID naturally capture this uncertainty... It is *natural* and *compact*.


The top-level MAID



The top-level MAID

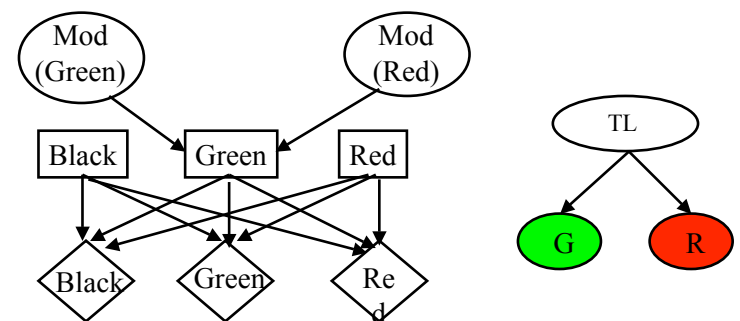


The top-level MAID




- At each block, each decision has a “Mod” parent defining the block used to make the decision.
- Probability distribution over “Mod” parent specifies the uncertainty of modeler.

Top level MAID




- When $\text{Mod}(\text{Green}) = G$, Black believes Green and Red use block G to make their decision.

Correlating between models



- “Mod” nodes can interact with other nodes in the MAID !
- When Collude=Green
 - Mod(Red)=G, Mod(Green) with some probability.

The top-level MAID



- Mod nodes can interact with other nodes in the MAID !

Specifying uncertainty over models

- At each block, each decision has a “Mod” parent defining the block used to make the decision.
- Probability dist. Over “Mod” parent specifies the uncertainty of modeler.
- “Mod” nodes can interact with other nodes in block MAID.