Equilibrium Refinements for Multi-Agent Influence Diagrams: Theory and Practice

PRESENTER: **James Fox**
james.fox@cs.ox.ac.uk

BACKGROUND:
Multi-agent influence diagrams (MAIDs) are a popular form of graphical model that can offer key advantages over traditional extensive form game (EFG) representations:

- **Compactness.**
- **Can exploit conditional independencies.**
- **Nash equilibria can be found more efficiently.**
- **A setting for causal queries (interventions or counterfactuals), which allows for incentives analysis in games.**

We introduce two important game theoretic concepts into MAIDs and prove several equivalence results between MAIDs and EFGs:

- **Subgames.**
- **Equilibrium refinements (Subgame Perfectness and Trembling Hand Perfectness).**

Key Definitions (see paper for formal definitions):

- A multi-agent influence diagram (MAID) is a triple (N, **V**, E) where N is a set of agents and (**V**, E) is a directed acyclic graph with **V** partitioned into decision, chance and utility variables for each agent. A multi-agent influence model (MAIM) is a parameterized MAID - it adds CPDs P(V|$\mathbf{Pa}_V$) to every non-decision variable $V \in \boldsymbol{V}/\boldsymbol{D}$.
- A Nash equilibrium (NE) is a policy profile π in a MAIM such that no agent may benefit by unilaterally deviating (i.e., $\forall\, i \in N, EU^i(\pi^i, \pi^{-i}) \geq EU^i(\hat{\pi}^i, \pi^{-i})\ \forall \hat{\pi}^i \in \Pi^i$)
- A node V is r-reachable from a decision $D \in \boldsymbol{D}^i$ in a MAID, M = (N, **V**, E) if a newly added parent $\hat{V}$ of V satisfies $\hat{V} \not\perp \boldsymbol{U}^i \cap \text{Desc}_D \mid \text{Fa}_D$.
- A MAID subgame is a subgraph of a MAID that is closed under r-reachability and mediators. A MAIM subgame is a parameterized MAID subgame with some setting of the variables outside the subgame.
- A full policy profile π is a MAIM-subgame perfect equilibrium in a MAIM M if π is an NE in every MAIM-subgame of M.

Multi-agent Influence Diagrams

offer a **compact** and **complete** **alternative graphical** representation of games

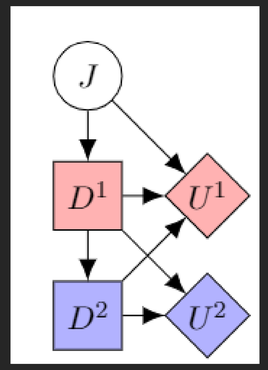

VS

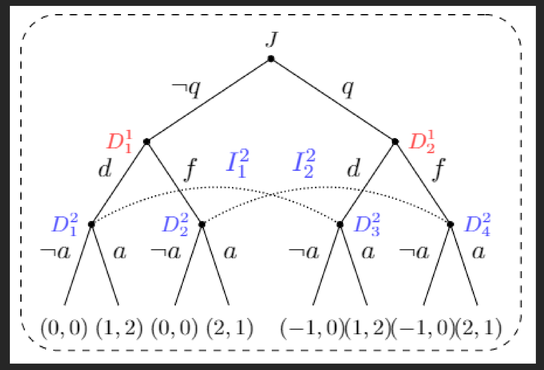

This work:

AAMAS 2021

Scan to **download** these **papers**

What's next?

Agent Incentives: A Causal Perspective (AAAI 2021)

Taxi Example:
A taxi-driver (agent 1) decides (D^1) whether to offer a discount to a potential customer (agent 2) depending on how many journeys it has already given that day (J). The customer decides (D^2) whether to accept a journey in this taxi depending on the price offered. The driver wants to maximise profit, which is greater if the customer buys at a higher price; however, if the driver hasn't given enough journeys that day and the customer rejects them, the driver must pay a penalty (the local council uses this penalty as a mechanism to prevent a proliferation of unnecessary taxis).

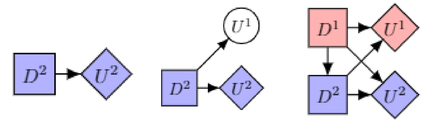

The three proper MAID subgames.

A subgame perfect equilibrium in a MAIM can rule our more non-credible threats.
Let D^1 ,D^2, and J be binary variables, J ~ Bernoulli(0.5), and $U^1 = D^1D^2 - J(1 - D^2)$ and $U^2 = D^2(3 - D^1)$ be the utility functions for the two agents. There are three pure NE, but only the latter is a MAIM SPE:

- **D^1=d always and D^2=a if and only if D^1=f**
- **D^1=f always and D^2=a if and only if D^1=f**
- D^1=f always and D^2=a always

D^2=¬a is always a non-credible threat because once the taxi has offered a price, the best response of the customer is to always accept the journey even at full price (if the customer observes D^1=f they will get expected utility 1 if they accept and 0 if they reject). This non-credible threat is ruled out by a subgame perfect equilibrium in a MAIM, but can't be ruled out in an EFG since there are no proper EFG subgames.

Work by James Fox*, Lewis Hammond*. Tom Everitt†, Alessandro Abate*, Michael Wooldridge*
* **University of Oxford**
† **DeepMind**

DeepMind

UNIVERSITY OF OXFORD