Improved Approximation of Interactive Dynamic Influence Diagrams Using Discriminative Model Updates

Prashant Doshi
University of Georgia, USA

Yifeng Zeng
Aalborg University, Denmark

Interactive Dynamic Influence Diagrams (I-DIDs)

I-DIDs are graphical models for decision making in multiagent settings. They are applicable to sequential decision making in extended interactions with other agents.

I-DIDs generalize dynamic IDs to multiagent domains and differ from MAIDs (Koller&Milch, UAI'01) and NIDs (Gal&Pfeffer, AAMAS'04).

I-DIDs offer online solutions to I-POMDPs (Gmytrasiewicz&Doshi, JAIR'05). They allow **nested modeling** of agents.

Related Work

We may **compact** the model space. Group together behaviorally equivalent models. Behaviorally equivalent models are those which result in identical solutions (policy trees).

Insight: Models spatially close (in belief) tend to be behaviorally equivalent. One approach is to use k-Means clustering and selection of models based on belief (Zeng et al., AAAI'07).

Problem: Many behaviorally equivalent models persist in the model space.

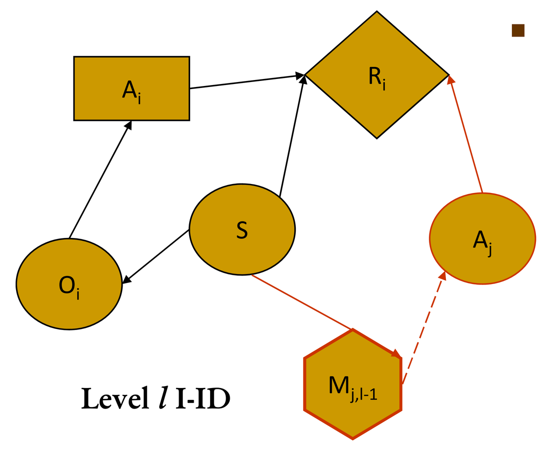

Level l I-ID

- A generic level l Interactive-ID (I-ID) for agent i situated with one other agent j
  - Model Node: $M_{j,l-1}$
    - Models of agent j at level l-1
  - Policy link: dashed arrow
    - Distribution over the other agent's actions given its models
  - Belief on $M_{j,l-1}$
    - $Pr(M_{j,l-1}|s)$
    - Update?

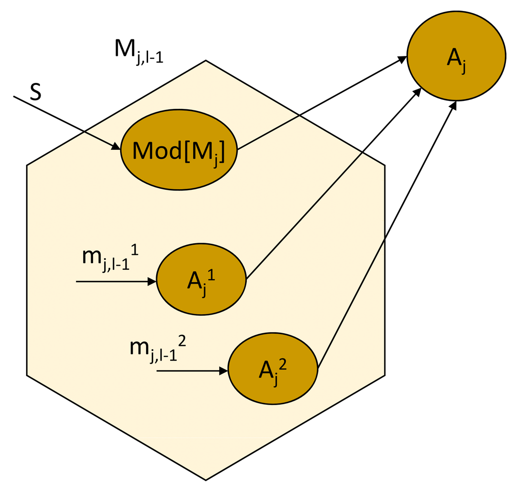

$m_{j,l-1}^1$, $m_{j,l-1}^2$ could be I-IDs or IDs

- Members of the model node
  - Different chance nodes are solutions of models $m_{j,l-1}$
  - Mod[M_j] represents the different models of agent j
- CPT of the chance node A_j is a multiplexer
  - Assumes the distribution of each of the action nodes (A_j^1, A_j^2) depending on the value of Mod[M_j]

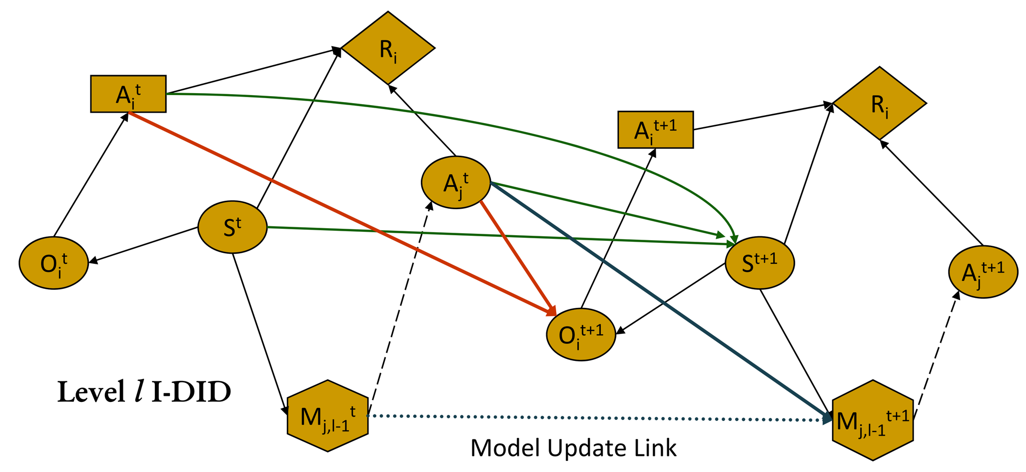

Level l I-DID

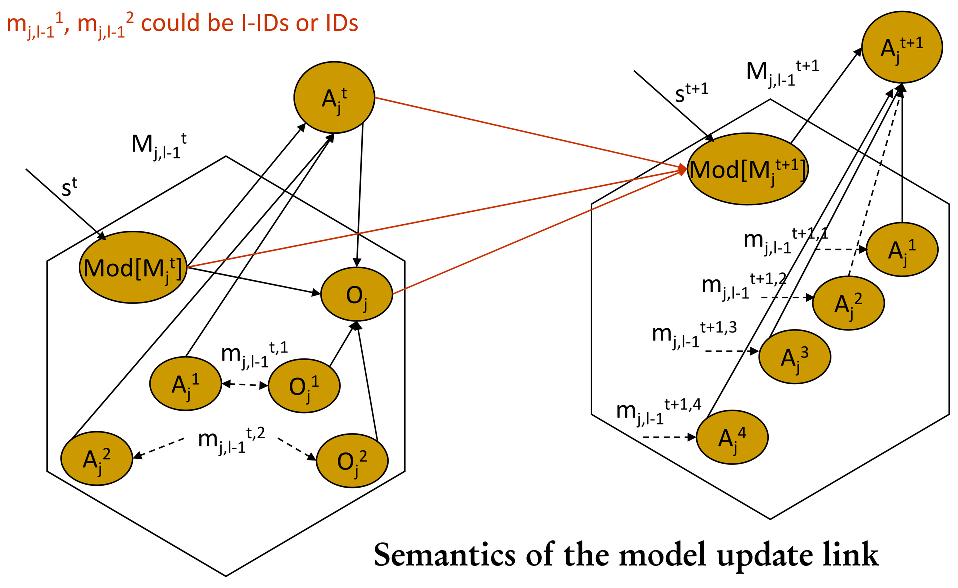

Semantics of the model update link

Discriminative Model Updates

Select only those models for updating which result in predictive behaviors distinct from others in updated model space. How do we identify these models?

We begin by merging individual policy trees bottom up to form a policy graph

Policy trees from individual models

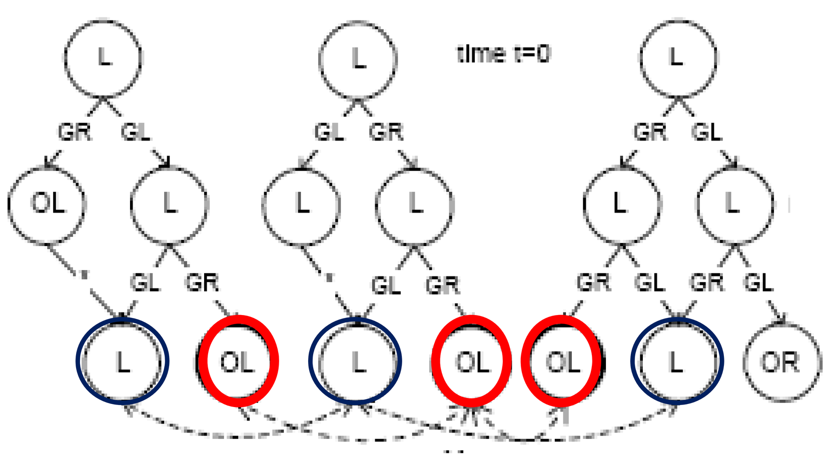

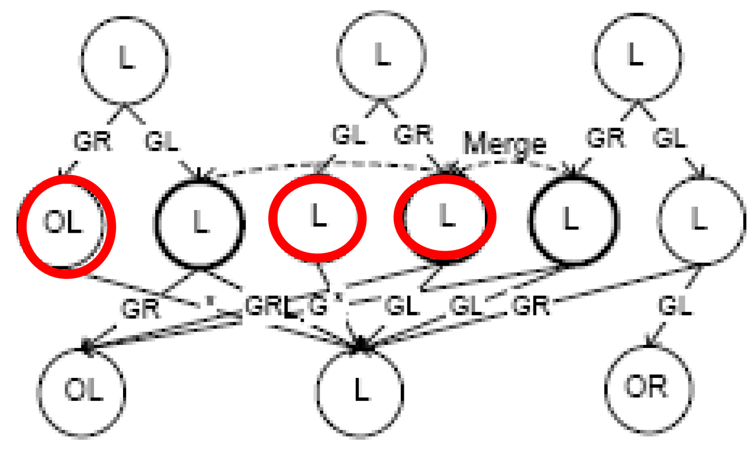

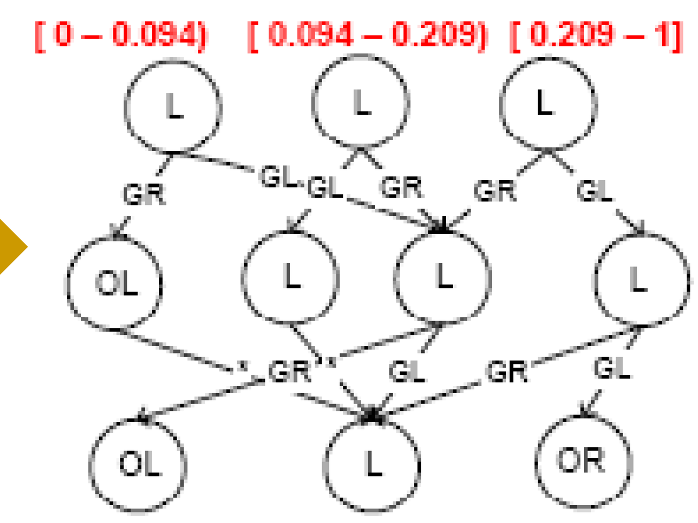

Consider 3 candidate models for agent j

$$\langle 0.01, \hat{\theta}_j \rangle, \langle 0.5, \hat{\theta}_j \rangle, \langle 0.05, \hat{\theta}_j \rangle$$

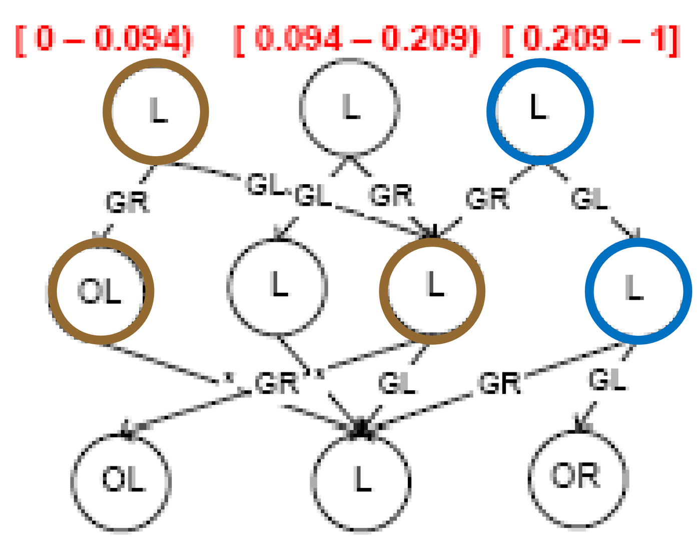

Model 1 will be updated using both L, GL and L, GR

Model 2 will be updated using L, GL but not L, GR because the latter update will result in a behaviorally equivalent model

Model 3 will not be updated at all as any updated model will be behaviorally equivalent

Approximation technique

We wish to avoid solving all models initially.

Approach

- Randomly select and solve K models recursively
- For the remaining models
  - If belief of a model is ε-close to solved model
    - Assume solution of solved model
  - otherwise, solve the model
- Merge solutions of solved models bottom up to obtain policy graph

Experimental Results

We compare our approach of using discriminative model updates (DMU) with a previous approximation technique: *k*-Means model clustering (MC).

We used two benchmark problems: Multiagent tiger and Multiagent machine maintenance.

I-DIDs are solved and the policies are executed in a simulated environment.

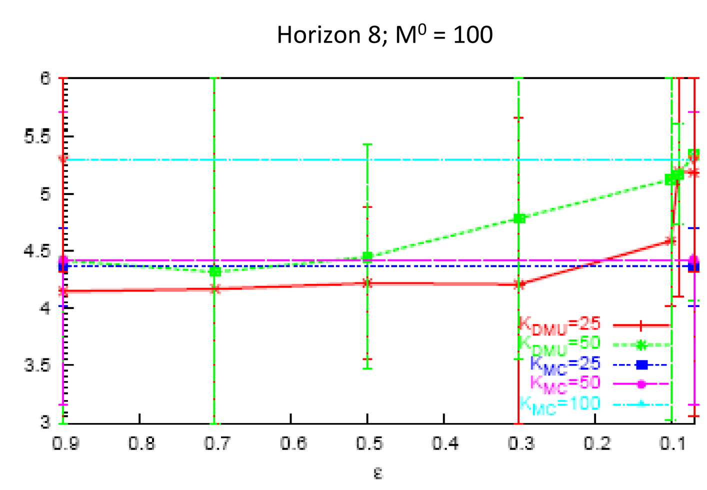

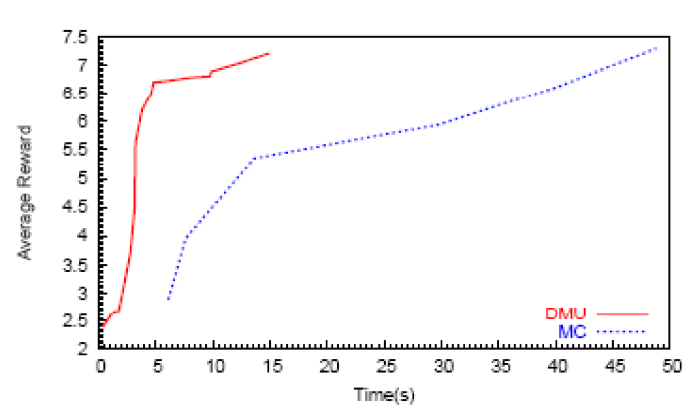

Multiagent tiger

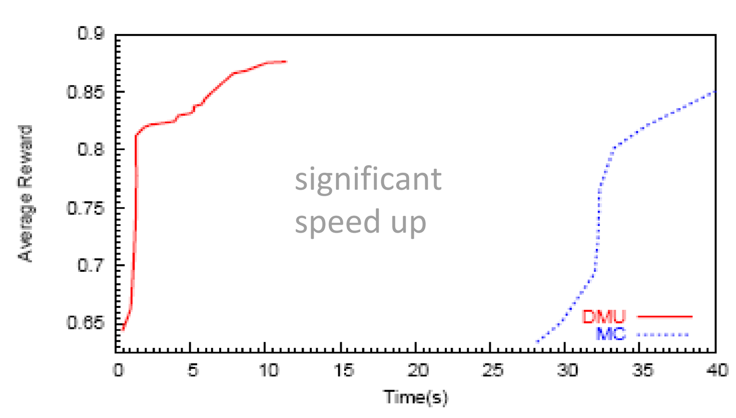

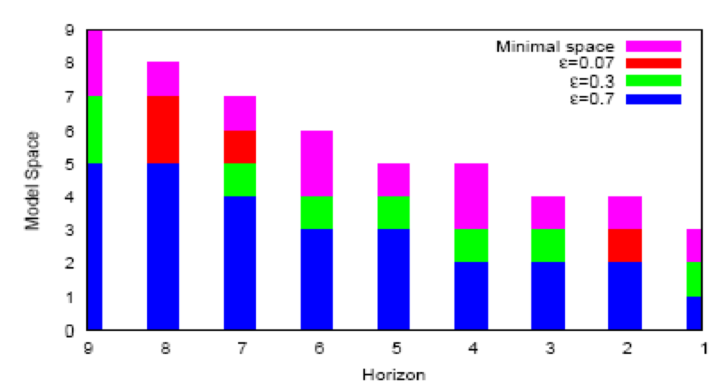

Multiagent machine maintenance

Discussion

I-DIDs offer a graphical formalism for sequential decision making in multiagent settings. They generalize DIDs to multiagent settings in a natural and intuitive manner.

We proposed an approach that preemptively avoids redundant updates. This leads to computational speed up. I-DIDs close to 20 horizons may be solved.

Level 1	T	Time (s)	
		DMU	MC
Tiger	6	2.53	19.86
	10	92.33	*
	17	488.12	*
MM	4	0.578	29.77
	10	95.31	*
	15	823.42	*

Acknowledgment

This research in supported in part by grant number #FA9550-08-1-0429 from the Air Force Office of Scientific Research (AFOSR).