# Advances in Multiagent Decision Making under Uncertainty 

Frans A. Oliehoek

Maastricht University

Coauthors: Matthijs Spaan (TUD), Shimon Whiteson (UvA), Nikos Vlassis (U.
Luxembourg), Jilles Dibangoye (INRIA), Chris Amato (MIT)

## Dynamics, Decisions \& Uncertainty

- Why care about formal decision making?


May 14, 2013

## Uncertainty

- Outcome Uncertainty

- Partial Observability

- Multiagent Systems: uncertainty about others


## Outline

- Background: sequential decision making
- Optimal Solutions of Decentralized POMDPs [JAIR'13]
- incremental clustering
- incremental expansion
- sufficient plan-time statistics [IJCAI'13]
- Other/current work
- Exploiting Structure [AAMAS'13]
- Multiagent RL under uncertainty [MSDM'13]


## Background: sequential decision making

## Single-Agent Decision Making

- Background: MDPs \& POMDPs
- An MDP $\left\langle S, A, P_{T}, R, h\right\rangle$
- $S$ - set of states

- $A$ - set of actions
- $P_{T}$ - transition function
$P\left(s^{\prime} \mid s, a\right)$
- $R$ - reward function
$R(s, a)$
- $h$ - horizon (finite)
- A POMDP $\left\langle S, A, P_{T}, O, P_{O}, R, h\right\rangle$
- $O$ - set of observations
- $P_{0}$ - observation function
$P\left(o \mid a, s^{\prime}\right)$


## Example: Predator-Prey Domain

- Predator-Prey domain
- 1 agent: predator
- prey is part of environment
- Formalization:
- states
- actions
- transitions
- rewards
$(-3,4)$


N,W,S,E
failing to move, prey moves
reward for capturing

## Example: Predator-Prey Domain

Markov decision process (MDP)

- Markovian state s... (which is observed!)
- policy $\pi$ maps states $\rightarrow$ actions


- Value function $\mathrm{Q}(\mathrm{s}, \mathrm{a})$
- Compute via value iteration / policy iteration

$$
Q(s, a)=R(s, a)+\gamma \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) \max _{a^{\prime}} Q\left(s^{\prime}, a^{\prime}\right)
$$

## Partial Observability

- Now: partial observability
- E.g., limited range of sight
- MDP + observations
- explicit observations
- observation probabilities
- noisy observations (detection probability)

$o=$ 'nothing '


## Partial Observability

- Now: partial observability
- E.g., limited range of sight
- MDP + observations
- explicit observations
- observation probabilities
- noisy observations (detection probability)

$o=(-1,1)$


## Partial Observability

- Now: partial observability
- E.g., limited range of sight
- MDP + observations
- explicit observations
- observation probabilities
- noisy observations (detection probability)

$o=(-1,1)$


## Can not observe the state

$\rightarrow$ Need to maintain a belief over states $b(s)$
$\rightarrow$ Policy maps beliefs to actions

$$
\pi(b)=a
$$

## Multiple Agents

- multiple agents, fully observable

Can coordinate based upon the state
$\rightarrow$ reduction to single agent: 'puppeteer' agent
$\rightarrow$ takes joint action

- Formalization:

- states
((3,-4), $(1,1),(-2,0))$
- actions
\{N,W,S,E\}
$\{(N, N, N),(N, N, W), \ldots,(E, E, E)\}$
- joint actions
- transitions
- rewards
probability of failing to move, prey moves
reward for capturing jointly


## Multiple Agents \& Partial Observability

- Dec-POMDP [Bernstein et al. '02]
- Reduction possible
$\rightarrow$ MPOMDP (multiagent POMDP)

- requires broadcasting observations!
- instantaneous, cost-free, noise-free communication $\rightarrow$ optimal [Pynadath and Tambe 2002]
- Without such communication: no easy reduction.


## Acting Based On Local Observations

- Acting on global information can be impractical:
" communication not possible
- significant cost (e.g battery power)
- not instantaneous or noise free

" scales poorly with number of agents!


## Formal Model

- A Dec-POMDP
- $\left\langle S, A, P_{T}, O, P_{O}, R, h\right\rangle$
- $n$ agents
- $S$ - set of states
- $A$ - set of joint actions
- $P_{T}$ - transition function
- O - set of joint observations
$o=\left\langle o_{1}, o_{2}, \ldots, o_{n}\right\rangle$
- $P_{0}$ - observation function
- $R$ - reward function
$P\left(o \mid a, s^{\prime}\right)$
$R(s, a)$
$a=\left\langle a_{1}, a_{2}, \ldots, a_{n}\right\rangle$
$P\left(s^{\prime} \mid s, a\right)$
- $h$ - horizon (finite)


## Running Example

- 2 generals problem
small army

large army



## Running Example

$S-\left\{\mathrm{S}_{\mathrm{L}}, \mathrm{S}_{\mathrm{S}}\right\}$
$A_{i}-\{(\mathrm{O})$ bserve, (A)ttack $\}$
$O_{i}-\{(\mathrm{L})$ arge, (S)mall $\}$
large army
Transitions

- Both Observe $\rightarrow$ no state change
- At least 1 Attack $\rightarrow$ reset ( $50 \%$ probability $\mathrm{S}_{\mathrm{L}^{\prime}} \mathrm{s}_{\mathrm{S}}$ )


## Observations

- Probability of correct observation: 0.85
- E.g., $P\left(<L, L>\mid S_{L}\right)=0.85 * 0.85=0.7225$

Rewards

- 1 general attacks $\rightarrow$ he loses the battle:
$R(*,<A, O>)=-10$
- Both generals Observe $\rightarrow$ small cost:
- Both Attack $\rightarrow$ depends on state:
$R(*,<O, O>)=-1$
$R\left(S_{L},<A, A>\right)=-20$
$R\left(S_{S^{\prime}}<A, A>\right)=+5$


## Off-line / On-line phases

- off-line planning, on-line execution is decentralized


## Planning Phase



Execution Phase


" (Smart generals make a plan in advance!)

## Goal of Planning

- Find an optimal joint policy

$$
\pi^{*}=\left\langle\pi_{1}, \pi_{2}\right\rangle \quad \pi_{i}: \vec{O}_{i} \rightarrow A_{i}
$$

- Value: expected sum of rewards:

$$
V(\pi)=\boldsymbol{E}\left[\sum_{t=0}^{h-1} R(s, a) \mid \pi, b^{0}\right]
$$

> No compact representation...

The problem is NEXP-complete [Bernstein et al. 2002]

- Also for $\varepsilon$-approximate solution! [Rabinovich et al. 2003]


## Should we give up on optimality?

- but we care about these problems...
- complexity: worst case
- may be able to optimally solve important problems
- optimal methods can provide insight in problems
- serve as inspiration for approximate methods
- need to benchmark: no usable upper bounds


## Advances in Exact Planning Methods

- Heuristic search + limitations
- Interpret search-tree nodes as 'Bayesian Games'
- Incremental Clustering
- Incremental Expansion
- Sufficient plan-time statistics


## Heuristic Search - 1

- Incrementally construct all (joint) policies
- 'forward in time'

1 joint policy


## Heuristic Search - 1

- Incrementally construct all (joint) policies
- 'forward in time'

1 partial joint policy

Start with unspecified policy

## Heuristic Search - 1

- Incrementally construct all (joint) policies
- 'forward in time'

1 partial joint policy


## Heuristic Search - 1

- Incrementally construct all (joint) policies
- 'forward in time'

1 partial joint policy


## Heuristic Search - 1

- Incrementally construct all (joint) policies
- 'forward in time'

1 complete joint policy (full-length)


## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 2

- Creating ALL joint policies $\rightarrow$ tree structure!



## Heuristic Search - 3

- too big to create completely...
- Idea: use heuristics
- avoid going down non-promising branches!
- Apply A* $\rightarrow$ Multiagent A* [Szer etal. 2005]


## Heuristic Search - 4

- Use heuristics $F(n)=G(n)+H(n)$
- $G(n)$ - actual reward of reaching $n$
- a node at depth t specifies $\varphi^{t}$ (i.e., actions for first t stages) $\rightarrow$ can compute $\mathrm{V}\left(\varphi^{\mathrm{t}}\right)$ over stages $0 . . . \mathrm{t}-1$
- H(n) - should overestimate!
- E.g., pretend that it is an MDP
- compute

$$
H(n)=H\left(\varphi^{t}\right)=\sum_{s} P\left(s \mid \varphi^{t}, b^{0}\right) \hat{V}_{M D P}(s)
$$

## Heuristics

- QPOMDP: Solve 'underlying POMDP'
- corresponds to immediate communication

$$
H\left(\varphi^{t}\right)=\sum_{\vec{\theta}^{\prime}} P\left(\vec{\theta}^{t} \mid \varphi^{t}, b^{0}\right) \hat{V}_{\text {POMDP }}\left(b^{\vec{b}^{\prime}}\right)
$$

- QBG corresponds to 1-step delayed communication
- Hierarchy of upper bounds [Oliehoek et al. 2008]

$$
Q^{*} \leq \hat{Q}_{k B G} \leq \hat{Q}_{B G} \leq \hat{Q}_{\text {POMDP }} \leq \hat{Q}_{M D P}
$$

## MAA* Limitations

- Number of children grows doubly exponentially with nodes depth
- For a node last stage, number of children: $O\left(\left|A_{*}\right|^{\left.\left.|n|\right|_{*}\right|^{\mid n-1}}\right)$
- Total number of joint policies:

$$
O\left(\left|A_{*}\right|^{\left.\left(n| |_{*} \mid-1\right)\right\rangle\left(\left|O_{*}\right|-1\right)}\right)
$$

$\rightarrow$ MAA* can only solve 1 horizon longer than brute force search... [Seuken \& Zilberstein '08]

- We introduce methods to fix this


## Collaborative Bayesian Games



- agents, actions
- types $\theta_{i} \leftrightarrow$ histories
- probabilities: $\mathrm{P}(\theta)$
- payoffs: Q( $\theta, a)$


## MAA* via Bayesian Games

- Each node $\leftrightarrow a \varphi^{t}$
- decision problem for stage t

|  | $\vec{\theta}_{2}^{t=0}$ | () |  |
| :---: | :---: | :---: | :---: |
| $\vec{\theta}_{1}^{t=0}$ |  | $a_{2}$ | $\bar{a}_{2}$ |
|  | $a_{1}$ | +2.75 | -4.1 |
| () | $\bar{a}_{1}$ | -0.9 | +0.3 |



## MAA* via Bayesian Games - 2

MAA* perspective


- node $\leftrightarrow \varphi^{t}$
- joint decision rule $\delta$ maps OHs to actions
- Expansion: appending all nextstage decision rules: $\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)$

BG perspective


- node $\leftrightarrow$ a BG
- joint BG policy $\beta$ maps 'types' to actions
- Expansion: enumeration of all joint BG policies $\varphi^{t+1}=\left(\varphi^{t}, \beta^{t}\right)$

```
direct correspondence: }\delta\leftrightarrow
```


## MAA* via Bayesian Games - 2



## The Decentralized Tiger Problem

- Two agents in a hallway
- States: tiger left ( $\mathrm{s}_{\mathrm{l}}$ ) or right ( $\mathrm{s}_{\mathrm{p}}$ )

- Actions: listen, open left, open right
- Observations: hear left (HL), hear right (HR)
- <Listen,Listen>
- 85\% prob. of getting right obs.
- e.g. $\left.\mathrm{P}(<\mathrm{HL}, \mathrm{HL}\rangle \mid<L i, L i>, S_{L}\right)=0.85 * 0.85=0.7225$
" otherwise: uniform random
- Reward: get the reward, acting jointly is better


## Lossless Clustering

- Two types (=action-observation histories) in a BG are probabilistically equivalent iff
$P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, a}\right)=P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, b}\right)$
$P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, a}\right)=P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, b}\right)$


(a) The joint type probabilities.

|  | $\vec{o}_{2}^{2}$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| $\vec{o}_{1}^{2}$ | $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | $\left(o_{\mathrm{HR}}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | 0.999 | 0.970 | 0.970 | 0.5 |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | 0.970 | 0.5 | 0.5 | 0.030 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HL}}\right)$ | 0.970 | 0.5 | 0.5 | 0.030 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ | 0.5 | 0.030 | 0.030 | 0.001 |

(b) The induced joint beliefs. Listed is the probability $\operatorname{Pr}\left(s_{l} \mid \overrightarrow{\boldsymbol{\theta}}^{2}, \boldsymbol{b}^{0}\right)$ of the tiger being behind the left door.

## Lossless Clustering

- Two types (=action-observation histories) in a BG are probabilistically equivalent iff
$P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, a}\right)=P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, b}\right)$
$P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, a}\right)=P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, b}\right)$

|  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| $\vec{o}_{1}^{2}$ | $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | $\left(o_{\mathrm{HR}}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | 0.261 | 0.047 | 0.047 | 0.016 |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | 0.047 | 0.016 | 0.016 | 0.047 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HL}}\right)$ | 0.047 | 0.016 | 0.016 | 0.047 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ | 0.016 | 0.047 | 0.047 | 0.261 |

(a) The joint type probabilities.

|  | $\vec{o}_{1}^{2}$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | $\left(o_{\mathrm{HR}}^{2}, o_{\mathrm{HL}}\right)$ | $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ |  |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HL}}\right)$ | 0.999 | 0.970 | 0.970 | 0.5 |
| $\left(o_{\mathrm{HL}}, o_{\mathrm{HR}}\right)$ | 0.970 | 0.5 | 0.5 | 0.030 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HL}}\right)$ | 0.970 | 0.5 | 0.5 | 0.030 |
| $\left(o_{\mathrm{HR}}, o_{\mathrm{HR}}\right)$ | 0.5 | 0.030 | 0.030 | 0.001 |

(b) The induced joint beliefs. Listed is the probability $\operatorname{Pr}\left(s_{l} \mid \overrightarrow{\boldsymbol{\theta}}^{2}, \boldsymbol{b}^{0}\right)$ of the tiger being behind the left door.

## Lossless Clustering

- Two types (=action-observation histories) in a BG are probabilistically equivalent iff
$P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, a}\right)=P\left(\vec{\theta}_{-i} \mid \vec{\theta}_{i, b}\right)$
$P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, a}\right)=P\left(s \mid \vec{\theta}_{-i}, \vec{\theta}_{i, b}\right)$
Clustering is lossless
restricting the policy space to clustered policies does not sacrifice optimality
- histories are bestresponse equivalent
- if criterion holds $\rightarrow$ same 'multiagent belief' $\mathrm{b}_{\mathrm{i}}\left(\mathrm{s}, \mathrm{q}_{-i}\right)$


## Incremental Clustering

- No need to cluster from scratch
- Probabilistic equivalence 'extends forwards'
- identical extensions of two PE histories are also PE
$\rightarrow$ can bootstrap from CBG of the previous stage
- 'Incremental clustering'


## Incremental Expansion

- Key idea: nodes have many children, but only few are useful.
- i.e., only few will be selected for further expansion
- others will have too low heuristic value

- if we can generate the nodes in decreasing heuristic order
$\rightarrow$ can avoid expansion of redundant nodes


## Incremental Expansion



## Incremental Expansion



Open list
a-7

## Incremental Expansion



## Incremental Expansion



## Incremental Expansion



## Incremental Expansion



## Incremental Expansion



## Incremental Expansion



## Incremental Expansion: How?

- How do we generate the next-best child?
- Node $\leftrightarrow$ BG, so...
- find the solutions of the BG

- in decreasing order of value
- i.e., 'incremental BG solver'
- Modification of BaGaBaB [Oliehoek et al. 2010]
- stop searching when next solution found
- save search tree for next time visited.
- Nested A*!


## Results

## GMAA*-ICE can solve higher horizons than listed

| $h$ | MILP | DP-LPC | DP-IPG | GMAA $-\mathrm{Q}_{\text {BG }}$ |  |  |  |
| :---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  | IC | ICE | heur |  |  |

## Results

| $h$ | $V^{*}$ | $T_{G M A A *}(\mathrm{~s})$ | $T_{I C}(\mathrm{~s})$ | $T_{I C E}(\mathrm{~s})$ |
| ---: | ---: | ---: | ---: | ---: |
| RECYCLING ROBOTS |  |  |  |  |
| 3 | 10.660125 | $\leq 0.01$ | $\leq 0.01$ | $\leq 0.01$ |
| 4 | 13.380000 | 713.41 | $\leq 0.01$ | $\leq 0.01$ |
| 5 | 16.486000 | - | $\leq 0.01$ | $\leq 0.01$ |
| 6 | $\mathbf{1 9 . 5 5 4 2 0 0}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 10 | $\mathbf{3 1 . 8 6 3 8 8 9}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 15 | $\mathbf{4 7 . 2 4 8 5 2 1}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 20 | $\mathbf{6 2 . 6 3 3 1 3 6}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 30 | $\mathbf{9 3 . 4 0 2 3 6 7}$ |  | 0.08 | 0.05 |
| 40 | $\mathbf{1 2 4 . 1 7 1 5 9 8}$ |  | 0.42 | 0.25 |
| 50 | $\mathbf{1 5 4 . 9 4 0 8 2 8}$ |  | 2.02 | 1.27 |
| 70 | $\mathbf{2 1 6 . 4 7 9 2 9 0}$ |  | - | 28.66 |
| 80 |  |  | - | - |

BROADCASTCHANNEL

| 4 | 3.890000 | $\leq 0.01$ | $\leq 0.01$ | $\leq 0.01$ |
| ---: | ---: | ---: | ---: | ---: |
| 5 | 4.790000 | 1.27 | $\leq 0.01$ | $\leq 0.01$ |
| 6 | $\mathbf{5 . 6 9 0 0 0 0}$ | - | $\leq 0.01$ | $\leq 0.01$ |
| 7 | $\mathbf{6 . 5 9 0 0 0 0}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 10 | $\mathbf{9 . 2 9 0 0 0 0}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 25 | $\mathbf{2 2 . 8 8 1 5 2 3}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 50 | $\mathbf{4 5 . 5 0 1 6 0 4}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 100 | $\mathbf{9 0 . 7 6 0 4 2 3}$ |  | $\leq 0.01$ | $\leq 0.01$ |
| 250 | $\mathbf{2 2 6 . 5 0 0 5 4 5}$ |  | 0.06 | 0.07 |
| 500 | $\mathbf{4 5 2 . 7 3 8 1 1 9}$ |  | 0.81 | 0.94 |
| 700 | $\mathbf{6 3 3 . 7 2 4 2 7 9}$ |  | 0.52 | 0.63 |
| 800 |  |  | - | - |
| 900 | $\mathbf{8 1 4 . 7 0 9 3 9 3}$ |  | 9.57 | 11.11 |
| 1000 |  |  | - | - |



GMAA*


GMAA*-ICE

Cases that compress well
May 14, 2013 * excluding heuristic

## Sufficient Plan-Time Statistics ${ }_{\text {[oieneoerer } 2013]}$

- Optimal decision rule depends on past joint policy $\varphi^{\dagger} \rightarrow$ search tree
- In fact possible to give an expression for the optimal value function based on $\varphi^{\dagger}$ [oliehoek etal. 2008]
- Recent insight: reformulation based on a sufficient statistic
- compact formulation of Q*
" search tree $\rightarrow$ DAG ("suff. stat-based pruning")


## Optimal Value Functions

2 parts:

- Value propagation:
- Value optimization:


## Optimal Value Functions

2 parts:

## (past Pol, AOH, decis. rule)

- Value propagation:
- last stage t=h-1

$$
Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)
$$

$$
\delta^{t}\left(\vec{\theta}^{t}\right)=\left\langle\delta_{1}^{t}\left(\vec{\theta}_{1}^{t}\right), \ldots, \delta_{n}^{t}\left(\vec{\theta}_{n}^{t}\right)\right\rangle
$$

- Value optimization:


## Optimal Value Functions

2 parts:

- Value propagation:
- last stage $\mathrm{t}=\mathrm{h}-1 \quad Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
- t<h-1

$$
\begin{array}{r}
Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{*+1}\right) \\
\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)
\end{array}
$$

- Value optimization:


## Optimal Value Functions

2 parts:

- Value propagation:
- last stage $\mathrm{t}=\mathrm{h}-1 \quad Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
" t<h-1
$Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{*+1}\right)$

$$
\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)
$$

- Value optimization:

$$
\delta^{*+1}=\operatorname{argmax}_{\delta^{t+1}} \sum_{\vec{\theta}^{t+1}} P\left(\vec{\theta}^{t+1} \mid b^{0}, \varphi^{t+1}\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

(need to do 'stage-wise' maximization)

## Optim

Qusv?

- state: $\varphi$

2 parts:

- actions: $\delta$

$$
V\left(\varphi^{t}\right)=\max _{\delta^{t}} Q^{*}\left(\varphi^{t}, \delta^{t}\right)
$$

- Value propagatio

$$
Q^{*}\left(\varphi^{t}, \delta^{t}\right)=\sum_{\vec{\theta}^{\prime}} P\left(\vec{\theta}^{t} \mid b^{0}, \varphi^{t}\right) Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)
$$

- last stage t=h-1
- t<h-1
$Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{* t+1}\right)$

$$
\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)
$$

- Value optimization:

$$
\delta^{*+1}=\operatorname{argmax}_{\mathrm{s}^{t+1}} \sum_{\vec{\theta}^{t+1}} P\left(\vec{\theta}^{t+1} \mid b^{0}, \varphi^{t+1}\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

(need to do 'stage-wise' maximization)

## Optimal Value Functions

2 parts:

- Value propagation:
- last stage $\mathrm{t}=\mathrm{h}-1 \quad Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
- t<h-1
$\left.Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{*+1}\right)\right)$
$\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)$
- Value optimization:

$$
\delta^{* t+1}=\arg \max _{\delta^{++1}} \sum_{\vec{\theta}^{t+1}} P\left(\vec{\theta}^{t+1} \mid b^{0}, \varphi^{t+1}\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

(need to do 'stage-wise' maximization)

## Optimal Value Functions

2 parts:

- Value propagation:
- last stage t=h-1 $\quad Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
- t<h-1
$\left.Q^{*}\left(\varphi^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+}, \delta^{*+1}\right)\right)$
$\varphi^{t+1}=\left(\varphi^{t}, \delta^{t}\right)$
- Value optimization:

$$
\delta^{*+1}=\operatorname{arg~max}_{\delta^{t+1}} \sum_{\vec{\theta}^{+1+}} P\left(\vec{\theta}^{t+1} \mid b\left(\varphi^{t+1}\right)\right) Q^{*}\left(\varphi^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

(need to do 'stage-wise' maximization)

## Optimal Value Functions

2 parts:

- Value propagation:
- last stage $\mathrm{t}=\mathrm{h}-1 \quad Q^{*}\left(\varphi^{h-1}, \vec{\theta}^{h-1}, \delta^{h-1}\right)=R\left(\vec{\theta}^{h-1}, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
- t<h-1

(need to do 'stage-wise' maximization)


## Optimal Value Functions

2 parts:

- Value propa\&
- last stage t $\left.\quad, \delta^{h-1}\left(\vec{\theta}^{h-1}\right)\right)$
- t<h-1


But: initial dependence only through this probability term!

- Value optimization:

(need to do 'stage-wise' maximization)


## Sufficient Statistic - 1

2 parts:

- Value propagation:

$$
Q^{*}\left(\sigma^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{*++1}\right)
$$

- Value optimization:

$$
\delta^{*+1}=\arg _{\max }^{\delta^{+1+1}} \sum_{\vec{\theta}^{t+1}} \sigma^{t+1}\left(\vec{\theta}^{t+1}\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

## Sufficient Statistic - 1

2 parts:

- Value propagation:

$$
Q^{*}\left(\sigma^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{*++1}\right)
$$

- Value optimization:

$$
\delta^{*+1}=\arg _{\max }^{\delta^{+1+1}} \sum_{\vec{\theta}^{t+1}} \sigma^{t+1}\left(\vec{\theta}^{t+1}\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

Limited use: every deterministic past joint policy induces a different $\sigma$ !

## Sufficient Statistic - 2

2 parts:
use: $\sigma^{t}\left(s, \vec{o}^{t}\right)$

- Value propagation:

$$
Q^{*}\left(\sigma^{t}, \vec{\theta}^{t}, \delta^{t}\right)=R\left(\vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right)+\sum_{o} P\left(o \mid \vec{\theta}^{t}, \delta^{t}\left(\vec{\theta}^{t}\right)\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{*++1}\right)
$$

- Value optimization:

$$
\delta^{*+1}=\arg _{\max }^{\delta^{+1+1}} \sum_{\vec{\theta}^{t+1}} \sigma^{t+1}\left(\vec{\theta}^{t+1}\right) Q^{*}\left(\sigma^{t+1}, \vec{\theta}^{t+1}, \delta^{t+1}\right)
$$

## Sufficient Statistic - 2

2 parts:
use: $\sigma^{t}\left(s, \vec{o}^{t}\right)$

- Value propagation:

- Value optimization:

$$
\left.\delta^{* t+1}=\arg m a x \delta_{\delta^{t+1}} \sum_{\vec{\theta}^{t+1}} \sigma^{t+}\left(\vec{\theta}^{t+1}\right) Q^{*}\left(\sigma^{t+1} \vec{\theta}^{t+1}\right) \delta^{t+1}\right)
$$

- substitute $\mathrm{AOH} \rightarrow \mathrm{OH}$
- but then $\rightarrow$ also adapt $\mathrm{R}(.$.$) and \mathrm{P}(\mathrm{o} \mid . .$.


## Sufficient Statistic - 2

2 parts:
use: $\sigma^{t}\left(s, \vec{o}^{t}\right)$

- Value propagation:

$$
Q^{*}\left(\sigma^{t}, \vec{o}^{t}, \delta^{t}\right)=R\left(\sigma^{t}, \vec{o}^{t}, \delta^{t}\right)+\sum_{o} P\left(o \mid \sigma^{t}, \vec{o}^{t}, \delta^{t}\right) Q^{*}\left(\sigma^{t+1}, \vec{o}^{t+1}, \delta^{*+1}\right)
$$

- Value optimization:

$$
\delta^{* t+1}=\operatorname{argmax}_{\delta^{t+1}} \sum_{\partial^{t+1}} \sigma^{t}\left(\partial^{t+1}\right) Q^{*}\left(\sigma^{t+1}, o^{t+1}, \delta^{t+1}\right)
$$

## Results -1

- Reduction in size of Q*

|  | $t=1$ |  | $t=2$ |  | $t=3$ |  |
| ---: | :--- | :--- | :--- | :--- | :--- | :--- |
|  | $\varphi_{1}$ | $\sigma_{1}$ | $\varphi_{2}$ | $\sigma_{2}$ | $\varphi_{3}$ | $\sigma_{3}$ |
| tiger | 9 | 2 | 729 | 20 | $4.78 e 6$ | 4520 |
| broadcast | 4 | 4 | 64 | 56 | $1.63 e 4$ | $1.16 e 4$ |
| recycling | 9 | 9 | 729 | 441 | $4.78 e 6$ | X |
| FF | 9 | 9 | 729 | 729 | $4.78 e 6$ | X |
| gridsmall | 25 | 16 | $1.56 e 4$ | 4096 | $6.10 e 9$ | X |
| hotel1 | 9 | 1 | $5.90 e 4$ | 4 | $1.7 e 19$ | - |

Table 1: Number of $\sigma_{t}$ vs. number of $\varphi_{t}$.

## Sufficient statistic-based pruning

- Before



## Sufficient statistic-based pruning

- Now
- many $\varphi \leftrightarrow$ same $\sigma$
- GMAA*-ICE with SSBP:
- perform GMAA*-ICE, but at each node compute $\sigma$
- if same $\sigma$ but lower G-value $\rightarrow$ prune branch


## Results - 2

- Speed-up GMAA*-ICE due to SSBP

|  | nodes created at depth $t$ |  |  |  |  |  |  |
| :---: | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  | SSBP | 1 | 2 | 3 | 4 | 5 | 6 |
| tiger |  |  |  |  |  |  |  |
| QMDP, h5 | yes | 1 | 10 | 615 | 28475 | 4 |  |
|  | no | 9 | 69 | 2319 | 41130 | 4 |  |
| QBG,h6 | yes | 1 | 2 | 8 | 18 | 162 | 1 |
| no | 9 | 2 | 8 | 18 | 166 | 1 |  |
| hotel1 |  |  |  |  |  |  |  |
| QMDP, h4 | yes | 1 | 4 | 6 | 3 |  |  |
|  | no | 9 | 252 | 11178 | 10935 |  |  |
| QMDP, h5 | yes | 1 | 4 | 12 | 15 | 7 |  |
|  | no | not solvable (out of 2 GB mem.) |  |  |  |  |  |
| QBG, h5 | no | 9 | 4 | 3 | 3 | 1 |  |

Table 2: Number of created child nodes in GMAA-ICE, when using sufficient statistic-based pruning (SSBP).


## References

- Most references can be found in

Frans A. Oliehoek. Decentralized POMDPs. In Wiering, Marco and van Otterlo, Martijn, editors, Reinforcement Learning: State of the Art, Adaptation, Learning, and Optimization, pp. 471-503, Springer Berlin Heidelberg, Berlin, Germany, 2012.

- Other:
" Dibangoye, Amato, Buffet, \& Charpillet. Optimally Solving Dec-POMDPs as ContinuousState MDPs. IJCAI, 2013.
- Oliehoek, Spaan, Amato, \& Whiteson. Incremental Clustering and Expansion for Faster Optimal Planning in Decentralized POMDPs. JAIR, 2013.
- Oliehoek. Sufficient Plan-Time Statistics for Decentralized POMDPs. IJCAI, 2013.

