

Reasoning about Knowledge for Planning

Bruno Zanuttini

Journée « Défis de l'IA », PFIA 2021, Tuesday, June 29, 2021



Partially Observable Planning Problems:

- ▶ states s including variables a, b (notation: $s \models a - b$)
- ▶ observations fr, ch
- ▶ obs. fr means $a \rightarrow a + 1$, ch means $b \rightarrow b + 1$

Sequences of observations:

- ▶ starting from $s \models 0 - 0$
- ▶ observing (ch, fr, fr, fr) means $s \models 3 - 1$
- ▶ but also (fr, ch, fr, fr) , (fr, fr, ch, fr) , (fr, fr, fr, ch)

Single-agent policies:

- ▶ assume single agent, $A = \{\text{didier}\}$
- ▶ assume action **att** is optimal when $s \models 1 - 0$
- ▶ assume action **def** is optimal when $s \models 3 - 1$
- ▶ standard policy representation (from $s \models 0 - 0$):
 - ▶ $\pi(\text{ch}, \text{fr}, \text{fr}, \text{fr}) = \text{def}$, $\pi(\text{fr}, \text{ch}, \text{fr}, \text{fr}) = \text{def}$, \dots , $\pi(\text{ch}) = \text{att}$
- ▶ knowledge-based representation: $\mathbf{K}(3 - 1) \mapsto \text{def}$, $\mathbf{K}(0 - 1) \mapsto \text{att}$

Multiagent setting:

- ▶ **collaborative** agents, (e.g., $A = \{\text{karim}, \text{kylian}, \text{n'golo}, \dots, \text{hugo}\}$)
- ▶ reason on own knowledge
- ▶ but also on **others'**: $\mathbf{K}(\neg \mathbf{K}_{\text{benjamin}}(\text{hugo-alone})) \mapsto \text{def}$

Some Planning Settings

Reasoning with Knowledge: The Single-Agent Case

Reasoning with Knowledge in QDec-POMDPs

More Challenges

Some Planning Settings

Reasoning with Knowledge: The Single-Agent Case

Reasoning with Knowledge in QDec-POMDPs

More Challenges

Planning actions for reaching a goal **contingent on observations**:

- ▶ state space, actions, goal states
- ▶ **known** successor function, **nondeterministic**
- ▶ **known** observation function

Example:


?	?	?
?	1	?
?	2	?
?	?	?

goal state = any with all non-bombs clicked

Stochastic setting:

Partially Observable Markov Decision Processes (POMDPs)

Mapping from histories of observations to actions



Execution: going from node to node using observations received

This presentation: **common rewards/collaborative planning**

Centralized setting:

- ▶ at each timestep, CA dictates each agent the action to take
- ▶ agent tell their observations to CA
- ▶ examples: sensor networks, patrolling

Distributed setting:

- ▶ no plan/policy agreed upon
- ▶ each agent plans, executes, observes **on its own**
- ▶ communication modelled by actions, **if any**
- ▶ example: crisis management (firemen, police...)

Still collaborative

Setting:

- ▶ each agent acts and observes on its own
- ▶ communication modelled by actions, if any
- ▶ but strategy/policy can be agreed upon beforehand

Examples:

- ▶ soccer
- ▶ exploration
- ▶ rescuing
- ▶ card games (Hanabi, Bridge...)...

Some Planning Settings

Reasoning with Knowledge: The Single-Agent Case

Reasoning with Knowledge in QDec-POMDPs

More Challenges

?	?	?
?	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

1	1	0
?	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

1	1	0
?	1	?
?	2	?
?	?	?

1	1	0
?	1	0
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

1	1	0
?	1	?
?	2	?
?	?	?

1	1	0
?	1	0
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

0	?	?
?	1	?
?	2	?
?	?	?

0	1	?
?	1	?
?	2	?
?	?	?

0	1	1
?	1	?
?	2	?
?	?	?

0	1	1
0	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

1	1	0
?	1	?
?	2	?
?	?	?

1	1	0
?	1	0
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

0	?	?
?	1	?
?	2	?
?	?	?

0	1	?
?	1	?
?	2	?
?	?	?

0	1	1
?	1	?
?	2	?
?	?	?

0	1	1
0	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

0	?	?
?	1	?
?	2	?
?	?	?

0	0	?
?	1	?
?	2	?
?	?	?

0	0	0
?	1	?
?	2	?
?	?	?

0	0	0
0	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

1	?	?
?	1	?
?	2	?
?	?	?

1	1	?
?	1	?
?	2	?
?	?	?

1	1	0
?	1	?
?	2	?
?	?	?

1	1	0
?	1	0
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

0	?	?
?	1	?
?	2	?
?	?	?

0	1	?
?	1	?
?	2	?
?	?	?

0	1	1
?	1	?
?	2	?
?	?	?

0	1	1
0	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?

0	?	?
?	1	?
?	2	?
?	?	?

0	0	?
?	1	?
?	2	?
?	?	?

0	0	0
?	1	?
?	2	?
?	?	?

0	0	0
0	1	?
?	2	?
?	?	?

?	?	?
?	1	?
?	2	?
?	?	?


0	?	?
?	1	?
?	2	?
?	?	?

0	0	?
?	1	?
?	2	?
?	?	?

0	0	0
?	1	?
?	2	?
?	?	?

0	0	0
1	1	?
?	2	?
?	?	?

Obvious, tree representation:



→ huge policies, huge search space

Knowledge state: all states of the problem consistent with observations

Observations:

0	0	0
1	1	?
?	2	?
?	?	?

Knowledge state: {

0	0	0
1	1	0
*	2	0
*	2	0

0	0	0
1	1	0
*	2	0
2	*	1

0	0	0
1	1	0
*	2	0
1	2	*

- ▶ MBP (Bertoli *et al.*, 2003):
represent knowledge states **compactly as OBDDs**
planning by forward search
- ▶ Contingent-FF (Hoffmann & Brafman, 2005):
check **equivalence of histories** at planning time
- ▶ Complexity of **belief tracking** = reasoning about histories
Bonet & Geffner, 2014; Brafman & Shani, 2016
- ▶ Finite-state controllers: policy as **transducer** observations/actions
- ▶ α -vectors for POMDPs:
 - ▶ assume reward of a is 10 in s_1 , -1 in s_2 , dual for b
 - ▶ **compact representation of policies** by $\vec{\alpha}(a) = (10, -1)$, $\vec{\alpha}(b) = (-1, 10)$
 - ▶ execute $\text{argmax}(\vec{\alpha}(a) \cdot \vec{b}_s)$

Pros:

- ▶ towards efficient planning algorithms
- ▶ ensure efficient execution

Cons:

- ▶ do not solve pb of succinct policies
- ▶ essentially reason belief-state-wise
- ▶ not readable

KBPs: specifying/verifying protocols (Fagin, Halpern, '90's)

Use for representing policies (with Lang, Saffidine, Schwarzentruher):

```
while  $\neg K \wedge_{i,j} (c_{i,j} \oplus m_{i,j})$  do  
  if  $K \neg m_{1,1}$  then click1,1 fi  
  if  $K \neg m_{1,2}$  then click1,2 fi  
  ...  
  if  $K \neg m_{H,W}$  then clickH,W fi  
od
```

Branching condition = compact representation of sets of KSs

Execution:

- ▶ maintain one's knowledge by progression through observations
- ▶ evaluate branching conditions in current knowledge

Drawback:

- ▶ computing what action to take: Θ_p^2 -hard (= $P^{NP}[\log n]$)
- ▶ tradeoff succinctness/execution
- ▶ suitable in not-so-real-time settings

Succinctness:

- ▶ always as succinct as standard representations
- ▶ possibly exponentially more so
- ▶ for some problems, than any reactive representation

Interpretability:

- ▶ readable, understandable
- ▶ “writeable:” let experts write policies, check auto.
- ▶ easier to write generic policies


Efficient execution:

- ▶ maintain agent's knowledge
explicit: OBDDs, implicit: SAT approaches
- ▶ evaluate branching conditions
- ▶ knowledge compilation questions

Efficient planning:

- ▶ very challenging
- ▶ forward search not so useful
- ▶ idea: start from policy tree and generalize
online regression (Rintanen, 2004)
- ▶ regression algorithms: need efficient languages
sets of sets of prop. assignments/sets of Boolean functions


Bienvenu, Fargier, Marquis 2010:



S5-DNF_{DNF,CNF} (BFM 2010):

$$V \left\{ \begin{array}{l} K(ab \vee cd) \wedge \neg K((x \vee y) \wedge (\bar{a})) \\ K(x) \wedge \neg K(z \vee \bar{x}) \wedge \neg K(a \vee b) \\ \neg K(t) \end{array} \right.$$

ESD (Niveau & Z. 2016):



Some Planning Settings

Reasoning with Knowledge: The Single-Agent Case

Reasoning with Knowledge in QDec-POMDPs

More Challenges

Reminder:

- ▶ execution is **decentralized**
- ▶ planning can be **centralized**
- ▶ policies: **per agent**, history \mapsto action

Alice in Toulouse, wants to join Bob in Paris. Can fly, but if pilot strike, will take train. Bob wants to pick her up at Paris airport or station.

Possible histories:

1. strike, web_ avail, web_ informs ...
2. strike, \neg web_ avail, \neg radio_ avail ...
3. strike, \neg web_ informs, \neg radio_ informs
4. \neg strike ...
5. etc.

Knowledge about the situation is not enough

- ▶ A knows strike, web_avail, web_informs
- ▶ must know that B browses the web
- ▶ agents must know each other's plan

Occupancy states (Dibangoye *et al.*, 2016):

- ▶ occupancy state = state + probability of joint histories
- ▶ sufficient planning-time statistics

Epistemic Planning (Bolander, Aucher...):

- ▶ actions with explicit epistemic effects
- ▶ typical example: broadcast/gossip
- ▶ essentially

Distributed setting (Engesser *et al.*, 2017):

- ▶ no centralized planning (crisis management)
- ▶ reason on what plans other can compute
- ▶ essentially studied in sequential setting

(With Lang, Saffidine, Schwarz)

- (A)
1. if \neg strike then fly else take train
 2. if train then
 - 2.1 if web_avail then browse web
 - 2.2 if radio_avail then listen to radio
 3. if $K_A K_B$ strike then go to station else go to airport
- (B)
1. wait
 2. do
 - 2.1 browse web
 - 2.2 listen to radio
 3. if $strike \wedge K_B K_A K_B$ strike then go to station else go to airport

Requires each agent to **evaluate epistemic formulas**

Essential assumption: each one **knows other agents' plans**

Possible histories:

1. $\text{strike, web_avail, web_informs} \dots : \mathbf{K_A K_B \text{strike}}$
2. $\text{strike, } \neg \text{web_avail, } \neg \text{radio_avail} \dots : \neg \mathbf{K_A K_B \text{strike}}$
3. $\text{strike, } \neg \text{web_informs, } \neg \text{radio_informs} : \mathbf{K_B } \neg \mathbf{K_A K_B \text{strike}}$
4. $\neg \text{strike} \dots : \mathbf{K_A } \neg \mathbf{K_B \text{strike}}$
5. etc.

Pros and cons:

- ▶ essentially same as single-agent setting
- ▶ hard execution (PSPACE-complete)
- ▶ exponential gains on succinctness

Representation languages for knowledge:

- ▶ same **knowledge compilation** questions
- ▶ space of multi-agent KS is **huge/infinite**
- ▶ use of compact languages:
through “programs”: Charrier
BDDs: Gattinger, Gamblin & Niveau

Some Planning Settings

Reasoning with Knowledge: The Single-Agent Case

Reasoning with Knowledge in QDec-POMDPs

More Challenges

Approximate reasoning:

- ▶ bound epistemic depth
- ▶ makes sense with bounded reasoners
- ▶ approx. valid plan...
- ▶ vs valid with approx. execution
- ▶ vs valid with approx. knowledge
- ▶ logics of belief

Learning:

- ▶ reinforcement learning (unknown model)
- ▶ build on PSRs (Littman)
- ▶ embedding into spaces of epistemic features