Epistemic Planning: Semantic Approach

Thomas Bolander, DTU Compute, Technical University of Denmark

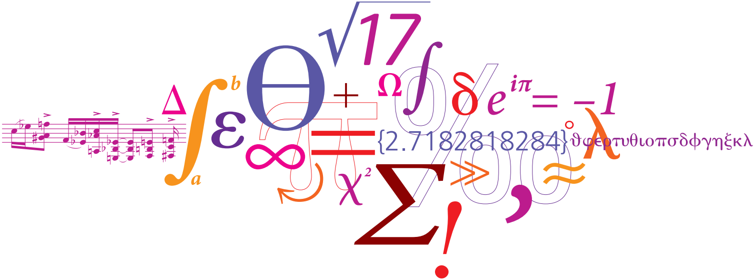

DTU Compute
Department of Applied Mathematics and Computer Science

Click link below to view video

http://www2.compute.dtu.dk/~tobo/children_cabinet_cropped.mp4

Epistemic planning $=$
automated *planning* $+$ Theory of Mind reasoning

Aim: To compute plans that can take the mental states of other agents into account.

Essentially: (Decentralised) **multi-agent planning** in environments with (potentially higher-order) **information asymmetry**.

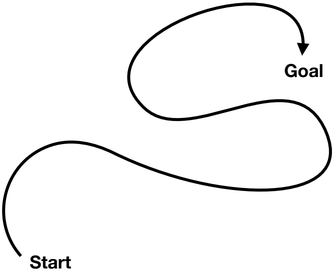

Automated planning

$+$

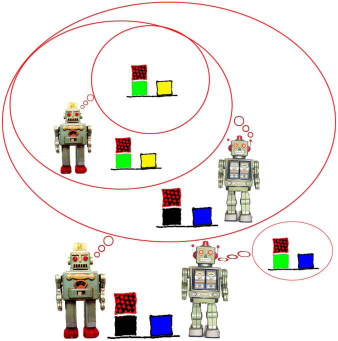

Logical reasoning about the mental states of other agents

Syntactic vs semantic, explicit vs implicit

When moving from standard propositional states to states including a Theory of Mind, there are two distinct paths one might take.

- **Syntactic approach**: States are (sets of) formulas (e.g. formulas of S5 epistemic logic)
- **Semantic approach**: States are semantic models (e.g. epistemic models = Kripke models).

Note: For propositional planning under full observability, the approaches are trivially equivalent.

Furthermore, for the semantic approach, there is a choice between:

- **Explicit approach**: Full state space is assumed given, and solution concept is defined directly in terms of this. E.g. logics like ATEL and CSL. [van der Hoek and Wooldridge, 2002, Jamroga and Aagotnes, 2007]
- **Implicit approach**: State space is induced by initial state and action library (as in classical STRIPS/PDDL planning).

DEL-based epistemic planning is *implicit* and *semantic*.

[Bolander and Andersen, 2011]

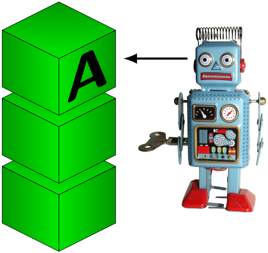

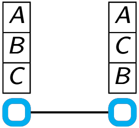

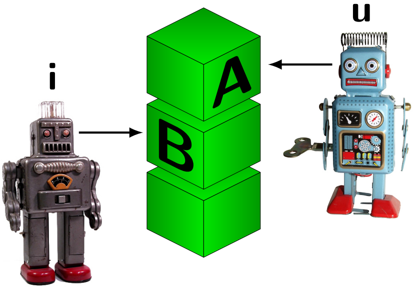

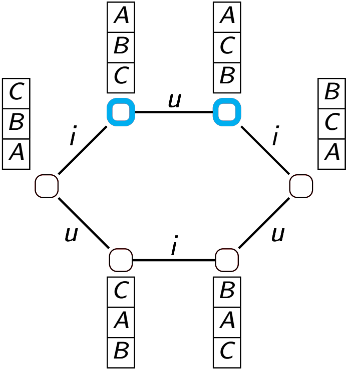

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**. **Designated worlds**: O (those considered possible by planning agent).

The coordinated attack problem in dynamic epistemic logic (DEL)

Two generals (agents), a and b. They want to coordinate an attack, and only win if they attack simultaneously.

d: "general a will attack at dawn".

m_i: the messenger is at general i (for $i = a, b$).

Initial **epistemic state**:

$s_0 =$

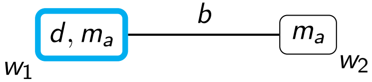

Nodes are **worlds**, edges are **indistinguishability edges** (reflexive loops not shown).

The coordinated attack problem in dynamic epistemic logic (DEL)

Recall: d means "a attacks at dawn"; m_i means messenger is at general i.

Available **epistemic actions** (aka **action models** aka **event models**):

$$a{:}send = \quad \underset{e_1}{\boxed{\begin{array}{ll} \text{pre}: & d \wedge m_a \\ \hline \text{post}: & m_b \wedge \neg m_a \end{array}}} \overset{a}{\text{———}} \underset{e_2}{\boxed{\begin{array}{ll} \text{pre}: & \top \\ \hline \text{post}: & \neg m_a \wedge \neg m_b \end{array}}}$$

And symmetrically an epistemic action $b{:}send$. We read $i{:}\alpha$ as "agent i does α".

Nodes are **events**, and each event has a **precondition** and a **postcondition** (effect). The precondition is an epistemic formula and the postcondition is a conjunction of literals.

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

The product update in dynamic epistemic logic

$s_0 =$

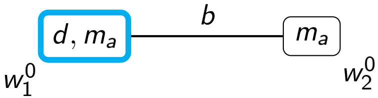

$s_0 \models K_a d \wedge \neg K_b d$

$a{:}send =$

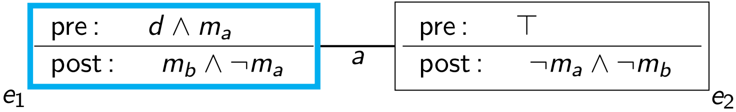

$s_0 \otimes a{:}send =$

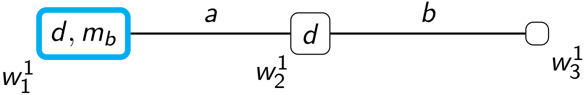

$s_0 \otimes a{:}send \models K_a d \wedge K_b d \wedge \neg K_a K_b d$

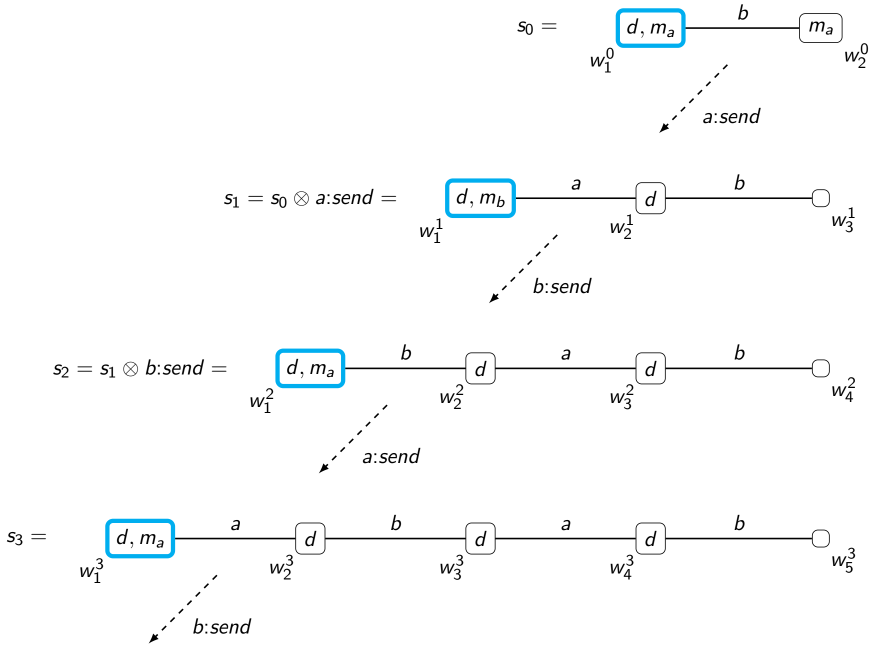

$s_0 =$
d, m_a w_1^0 — b — m_a w_2^0
$a{:}send$
$s_1 = s_0 \otimes a{:}send =$
d, m_b w_1^1 — a — d w_2^1 — b — w_3^1
$b{:}send$
$s_2 = s_1 \otimes b{:}send =$
d, m_a w_1^2 — b — d w_2^2 — a — d w_3^2 — b — w_4^2
$a{:}send$
$s_3 =$
d, m_a w_1^3 — a — d w_2^3 — b — d w_3^3 — a — d w_4^3 — b — w_5^3
$b{:}send$

Epistemic planning tasks

Definition. An **epistemic planning task** (or simply a **planning task**) $T = (s_0, A, \gamma)$ consists of an epistemic state s_0 called the **initial state**; a finite set of epistemic actions A; and a **goal formula** γ of the epistemic language.

Definition. A (sequential) **solution** to a planning task $T = (s_0, A, \gamma)$ is a sequence of actions $\alpha_1, \alpha_2, \dots, \alpha_n$ from A such that for all $1 \leq i \leq n$, α_i is applicable in $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_{i-1}$ and

$$s_0 \otimes \alpha_1 \otimes \alpha_2 \otimes \cdots \otimes \alpha_n \models \gamma.$$

Example. Let s_0 be the initial state of the coordinated attack problem. Let $A = \{a{:}send, b{:}send\}$. Then the following are planning tasks:

1. $T = (s_0, A, Cd)$, where C denotes common knowledge. It has no solution.
2. $T = (s_0, A, E^n d)$, where E denotes "everybody knows" and $n \geq 1$. It has a solution of length n.

[Bolander et al., 2020]

Epistemic planning example: Get the cube

- **Objects**: $\mathcal{O} = \{b_1, b_2, c\}$, two boxes b_1 and b_2, and a cube c.
- **Agents**: $\mathcal{A} = \{h, a\}$, a human h and a robot r. The robot is the planning agent.
- **Atomic propositions**: $In(x, y)$ means x is in y, where $x, y \in \mathcal{O} \cup \mathcal{A}$ (when $y \in \mathcal{A}$, it means y is holding x).

Initial epistemic state:

$$s_0 = \quad \boxed{In(c, b_1)} \overset{h}{\rule[0.5ex]{4em}{0.4pt}} \boxed{In(c, b_2)}$$

The goal is for the human to hold the red cube, $In(r, h)$.

Actions specialised for the case of $\mathcal{O} = \{b_1, b_2, c\}$.

Agent i (semi-privately) **peeks** into box x:

$$i{:}peek(x) = \quad \boxed{\text{pre}: \quad In(c,x)} \overset{}{\underset{\mathcal{A}-\{i\}}{\text{———}}} \boxed{\text{pre}: \quad \neg In(c,x)}$$

Agent i (publicly) **picks up** object x from y:

$$i{:}pickup(x,y) = \quad \boxed{\begin{array}{ll} \text{pre}: & In(x,y) \\ \hline \text{post}: & In(x,i) \land \neg In(x,y) \end{array}}$$

Agent i (publicly) **puts** object x in y:

$$i{:}putdown(x,y) = \quad \boxed{\begin{array}{ll} \text{pre}: & In(x,i) \\ \hline \text{post}: & In(x,y) \land \neg In(x,i) \end{array}}$$

Agent i (publicly) **announces** that formula φ is true:

$$i{:}ann(\varphi) = \quad \boxed{\text{pre}: \quad \varphi}$$

Get the cube: Planning task and solutions

The planning task T has the actions of the previous slide and initial state s_0 and goal γ given by:

$$s_0 = \quad \boxed{In(c, b_1)} \overset{h}{\text{———}} \boxed{In(c, b_2)} \qquad\qquad \gamma = In(r, h)$$

Solution to T, by robot R:

$$\begin{array}{rl}
s_0 = & \boxed{In(c, b_1)} \overset{h}{\text{———}} \boxed{In(c, b_2)} \\
& \quad \Big\downarrow \; r{:}pickup(c, b_1) \\
s_1 = s_0 \otimes r{:}pickup(c, b_1) = & \boxed{In(c, r)} \\
& \quad \Big\downarrow \; r{:}putdown(c, h) \\
s_2 = s_1 \otimes r{:}putdown(c, h) = & \boxed{In(c, h)}
\end{array}$$

Applicability, perspective shifts, implicit coordination

Seemingly simpler solution: $h{:}pickup(c, b_1)$. But intuitively, this shouldn't work, since the human doesn't know the cube is in box 1...

Applicability: An action α is **applicable** in a state s if for each designated world w of s there is a designated event e of α with $w \models pre(e)$.

Perspective shift: The **perspective shift** of state s to agent i, denoted s^i, is achieved by closing under the indistinguishability relation of i. We call s^i the **perspective** of agent i on state s.

$$s_0 = \boxed{In(c, b_1)} \overset{h}{\text{———}} \boxed{In(c, b_2)} \qquad s_0^h = \boxed{In(c, b_1)} \overset{h}{\text{———}} \boxed{In(c, b_2)}$$

Example. $h{:}pickup(c, b_1)$ is not applicable in s_0 from h's perspective.

Implicitly coordinated solution to planning task: Each action has to be applicable from the perspective of the acting agent; and the product update $s \otimes i{:}\alpha$ is replaced by $s^i \otimes i{:}\alpha$.

Get the cube: Implicit coordination

Joint solution to T, by robot R, implicitly coordinated:

$$s_0 = \quad \boxed{In(c, b_1)} \overset{h}{\rule[0.5ex]{3em}{0.4pt}} \boxed{In(c, b_2)}$$

$$\Big\downarrow \; r{:}ann(In(c, b_1))$$

$$s_1 = s_0 \otimes r{:}ann(In(c, b_1)) = \quad \boxed{In(c, b_1)}$$

$$\Big\downarrow \; h{:}pickup(c, b_1)$$

$$s_2 = s_1 \otimes h{:}pickup(c, b_1) = \quad \boxed{In(c, h)}$$

If purely epistemic actions (announcements) have a lower cost than ontic actions (moving things around), the solution above is the only optimal one.

Undecidability: lengthening and shortening chains

Consider a chain produced by the coordinated attack problem:

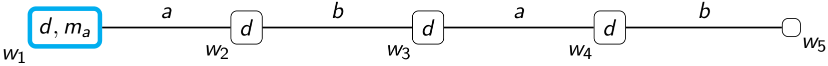

Using preconditions of modal depth 1 we can shorten the chain by 1:

shorten =

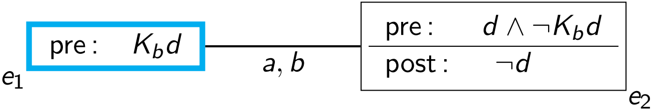

We can now both lengthen (by *send*) and shorten chains (by *shorten*), and this allows us to encode two-counter machines $\Rightarrow$ undecidability of the plan existence problem!

Undecidability holds even with preconditions of modal depth 1, and for purely epistemic planning (no postconditions) even for modal depth 2. [Bolander and Andersen, 2011, Charrier et al., 2016, Bolander et al., 2020]

Some of the current challenges in epistemic planning

- **Undecidability issues**: open complexity problems.
  [Bolander et al., 2020]
- **State size explosion problems**: find compact state representations.
  [Charrier and Schwarzentruber, 2017, van Benthem et al., 2018]
- **The belief-revision problem in DEL**: How to recover from false beliefs without an underlying epistemic relation. Relates to the state size explosion problem.
- **Heuristics for epistemic planning**: to reduce all of the above mentioned complexity and scalability issues
- **Languages**: syntactic languages for describing actions.
  [Baral et al., 2012, Baral et al., 2013]

This, and much more, is discussed in the "Epistemic Planning" special issue of AIJ currently being finalised.

References I

Baltag, A., Moss, L. S. and Solecki, S. (1998).
The Logic of Public Announcements and Common Knowledge and Private Suspicions.
In Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98), (Gilboa, I., ed.), pp. 43–56, Morgan Kaufmann.

Baral, C., Gelfond, G., Pontelli, E. and Son, T. C. (2012).
An action language for reasoning about beliefs in multi-agent domains.
In Proceedings of the 14th International Workshop on Non-Monotonic Reasoning vol. 4,.

Baral, C., Gelfond, G., Pontelli, E. and Son, T. C. (2013).
Reasoning about the beliefs of agents in multi-agent domains in the presence of state constraints: The action language mal.
In International Workshop on Computational Logic in Multi-Agent Systems pp. 290–306, Springer.

Bolander, T. and Andersen, M. B. (2011).
Epistemic Planning for Single- and Multi-Agent Systems.
Journal of Applied Non-Classical Logics *21*, 9–34.

Bolander, T., Charrier, T., Pinchinat, S. and Schwarzentruber, F. (2020).
DEL-based Epistemic Planning: Decidability and Complexity.
Artificial Intelligence *287*, 1–34.

Charrier, T., Maubert, B. and Schwarzentruber, F. (2016).
On the Impact of Modal Depth in Epistemic Planning.
In Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, USA, July 12-15, 2016.

Charrier, T. and Schwarzentruber, F. (2017).
A Succinct Language for Dynamic Epistemic Logic.
In AAMAS pp. 123–131,.

References II

Jamroga, W. and Aagotnes, T. (2007).
Constructive knowledge: what agents can achieve under imperfect information.
Journal of Applied Non-Classical Logics *17*, 423–475.

van Benthem, J., van Eijck, J., Gattinger, M. and Su, K. (2018).
Symbolic Model Checking for Dynamic Epistemic Logic — S5 and Beyond.
Journal of Logic and Computation *28*.

van der Hoek, W. and Wooldridge, M. (2002).
Tractable Multiagent Planning for Epistemic Goals.
In In Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002) pp. 1167–1174, ACM Press.

van Ditmarsch, H. and Kooi, B. (2008).
Semantic Results for Ontic and Epistemic Change.
In Logic and the Foundation of Game and Decision Theory (LOFT 7), (Bonanno, G., van der Hoek, W. and Wooldridge, M., eds), Texts in Logic and Games 3 pp. 87–117, Amsterdam University Press.