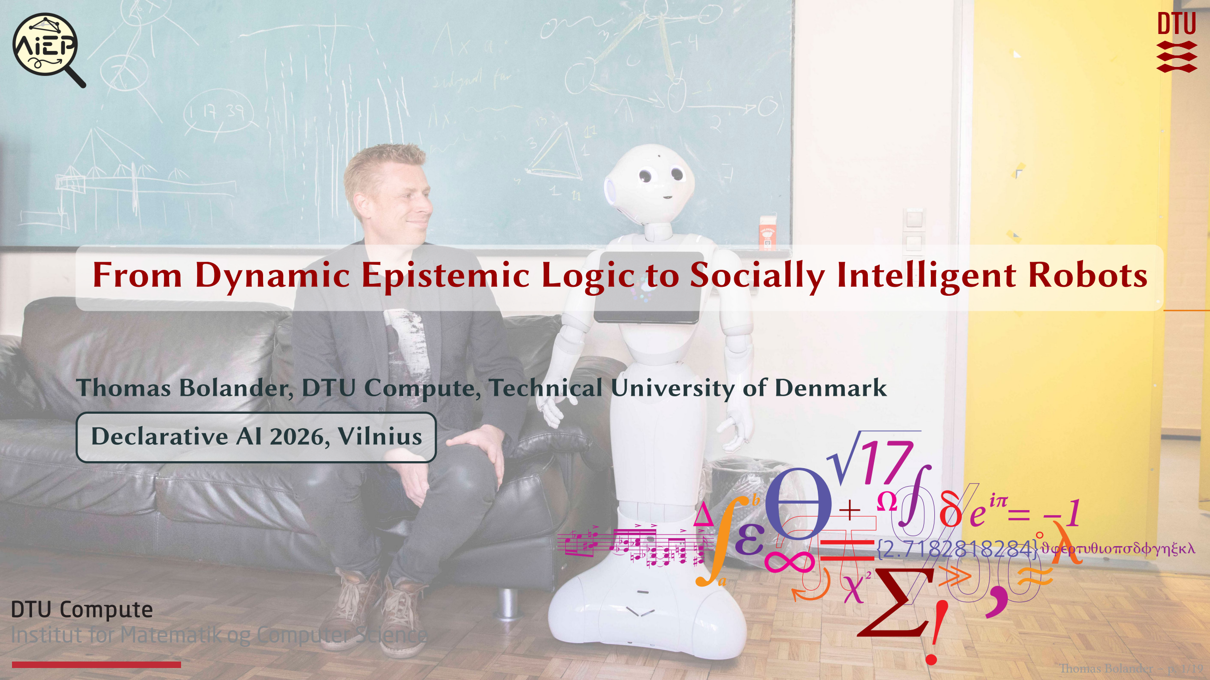

From Dynamic Epistemic Logic to Socially Intelligent Robots

Thomas Bolander, DTU Compute, Technical University of Denmark

Declarative AI 2026, Vilnius

Automated planning: Classical planning tasks

Definition (Adapted from [Ghallab et al., 2004])

A **planning task** is $T = (s_0, \mathcal{A}, \varphi_g)$, where

- s_0 is an **initial state**.
- $\mathcal{A}$ is a finite set of available **actions**. $\mathcal{A}$ induces a transition function $\circ : (s, \alpha) \mapsto s'$ where $s' = s \circ \alpha$ is the state resulting from applying action $\alpha \in \mathcal{A}$ in state s.
- φ_g is a **goal formula**.

A **solution** to a planning task $(s_0, \mathcal{A}, \varphi_g)$ is a sequence of actions (a **plan**) $\pi = \alpha_1, \ldots, \alpha_n$ from $\mathcal{A}$ such that $s \circ \alpha_1 \circ \cdots \circ \alpha_n \models \varphi_g$.

State space of a task T:

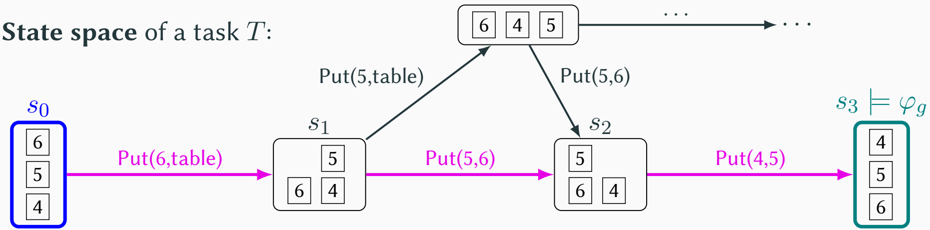

Propositional planning: Propositional states and action schemas

Action schema describing the $Put(x,y)$ action for putting object x on top of object y:

$$\begin{array}{|l|}\hline \textsc{Action}: Put(x,y) \\ \textsc{Precondition}: On(x,z) \wedge \cdots \\ \textsc{Effect}: On(x,y) \wedge \neg On(x,z) \\ \hline \end{array}$$

In STRIPS/PDDL

$$\begin{array}{|ll|}\hline \text{pre}: & On(x,z) \wedge \cdots \\ \hline \text{post}: & \begin{array}{l} On(x,y) := \top \\ On(x,z) := \bot \end{array} \\ \hline \end{array}$$

In Dynamic Epistemic Logic (DEL)

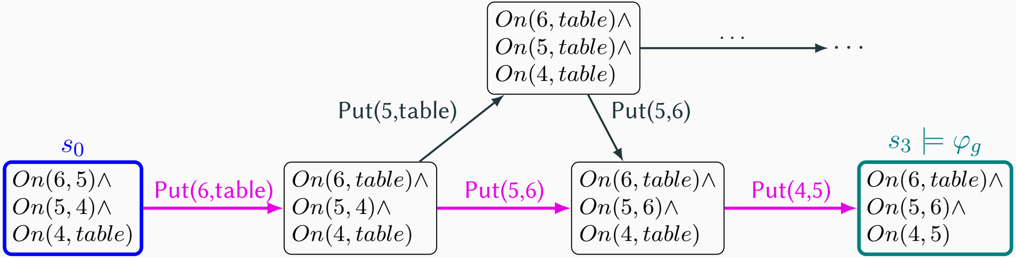

Link to movie (click the link):
http://www2.compute.dtu.dk/~tobo/children_cabinet_cropped.mov

Epistemic states: (Multipointed) S5 Kripke models

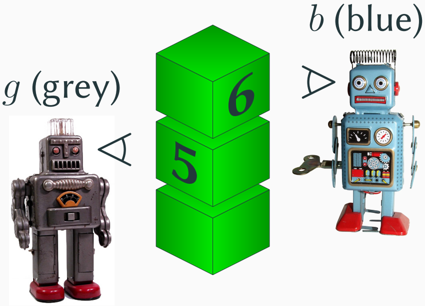

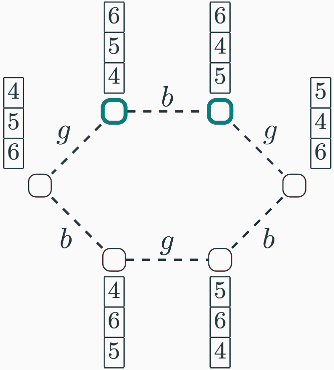

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.

Designated worlds: ▢ (those considered possible by modelling agent, b in this case).

$s \models \neg K_b \mathit{middle}(5) \wedge \neg K_b \mathit{middle}(4) \wedge K_b(K_g \mathit{middle}(5) \vee K_g \mathit{middle}(4))$

Agent b to agent g: "Which letter does the middle block have?"

Dynamic epistemic logic (DEL) by example: The product update operator $\otimes$

$s_0 =$

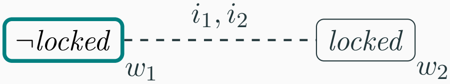

i_1, i_2, i_3 are agents

$i_1{:}open =$

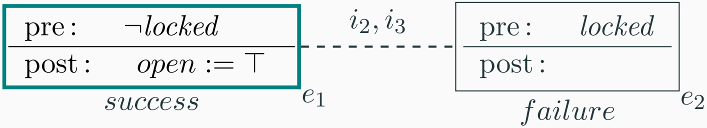

$s_0 \otimes i_1{:}open =$

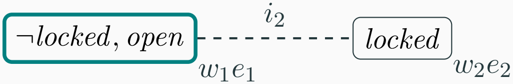

$\otimes$ is the transition function

$i_1{:}open$ is an **event model** (**epistemic action**) representing "agent i_1 does *open*". Nodes are **events**. Each event has a **precondition** and **postconditions** for all atoms (all of which are epistemic formulas).

Designated events of action $i{:}\alpha$: ▢ (those considered possible by i, closed under i's indistinguishability relation).

Planning based on DEL: Epistemic planning tasks

Definition

An **(epistemic) planning task** $T = (s_0, \mathcal{A}, \varphi_g)$ consists of

- **Initial state** s_0: An epistemic state (a multipointed Kripke model).
- **Actions** $\mathcal{A}$: Finite set of epistemic actions (multipointed event models).
- **Goal formula** φ_g: formula of epistemic logic.

Definition

A (sequential) **solution** to a planning task $T = (s_0, \mathcal{A}, \varphi_g)$ is an action sequence (a **plan**) $i_1{:}\alpha_1, i_2{:}\alpha_2, \ldots, i_n{:}\alpha_n \in \mathcal{A}$ such that

$$s_0 \otimes (i_1{:}\alpha_1) \otimes (i_2{:}\alpha_2) \otimes \cdots \otimes (i_n{:}\alpha_n) \models \varphi_g.$$

Defining $(\![\alpha]\!)\varphi := \langle\alpha\rangle\top \wedge [\alpha]\varphi$, this can be reformulated as

$$s_0 \models (\![i_1{:}\alpha_1]\!)(\![i_2{:}\alpha_2]\!)\cdots(\![i_n{:}\alpha_n]\!)\varphi_g.$$

Epistemic planning task example: Human-robot collaboration

Epistemic state s_0

$$\boxed{In(cube, box_1)} \overset{r,h}{\text{- - - - - - - -}} \boxed{In(cube, box_2)}$$

$$\otimes$$

Epistemic action $r{:}perceive(In(cube, box_1))$

pre :	$In(cube, box_1)$	h	pre :	$\neg In(cube, box_1)$
post :			post :	

$$=$$

Resulting state $s_0 \otimes r{:}perceive(In(cube, box_1))$

$$\boxed{In(cube, box_1)} \overset{h}{\text{- - - - - - - -}} \boxed{In(cube, box_2)}$$

Goal formula:

$\varphi_g = \bigvee_{x \text{ is box}} K_h In(cube, x)$

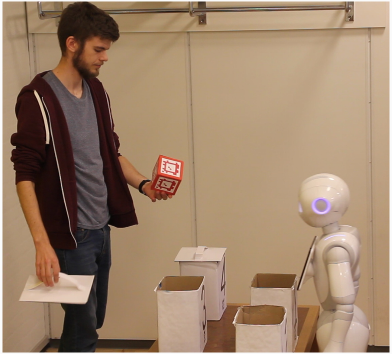

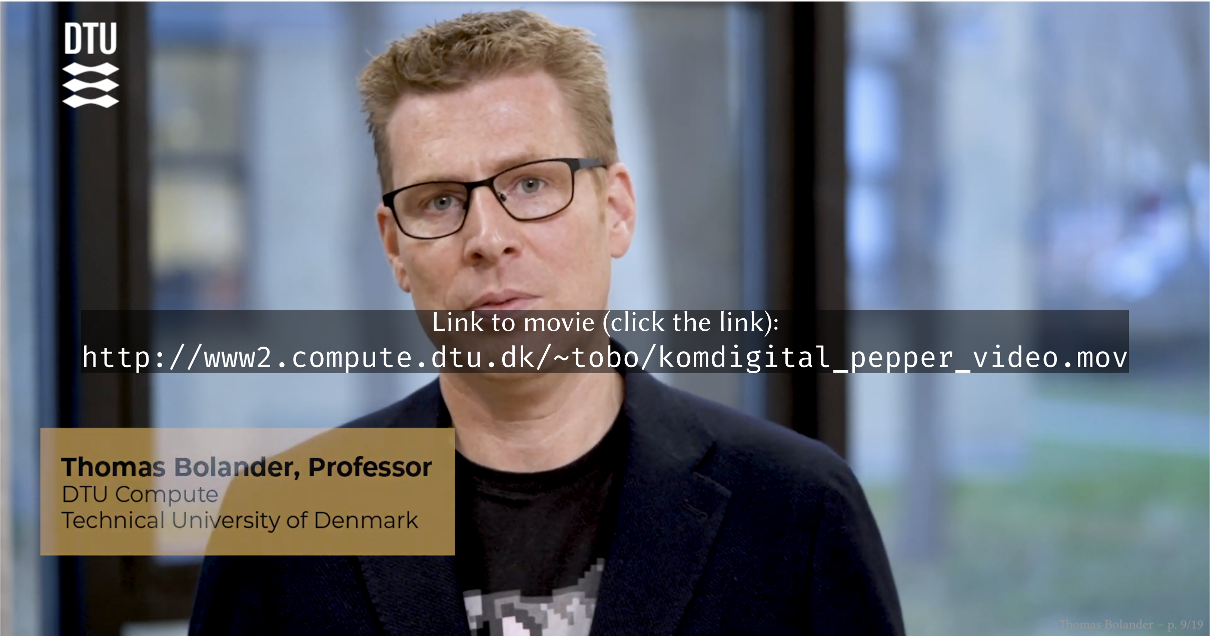

Link to movie (click the link):

http://www2.compute.dtu.dk/~tobo/komdigital_pepper_video.mov

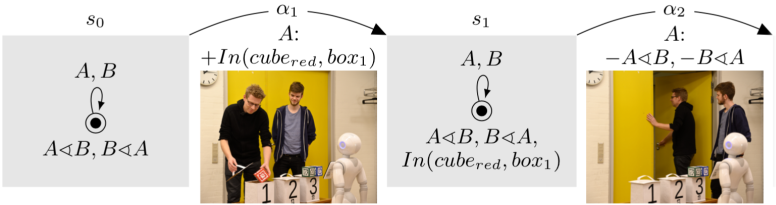

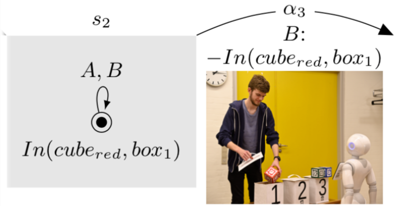

Atom $In(c_1, c_2)$: obj. c_1 is in container c_2.

Atom $a \lhd b$ (a sees b): agent a currently observes the actions of agent b.

Action $i{:}X$: X is a list of assignments ($+p$ and $-p$ where p is an atom).

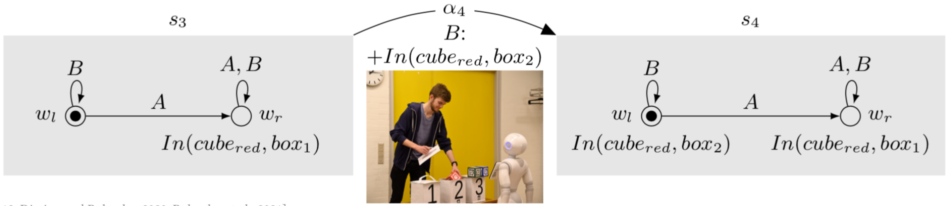

Perception Layer:
deep neural networks
track objects, agents,
actions (subsymbolic)

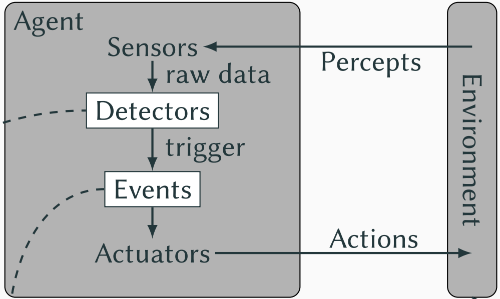

Cognition Layer:
DEL + planning (symbolic)

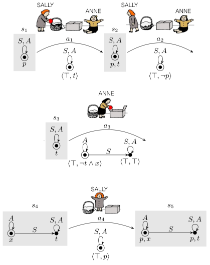

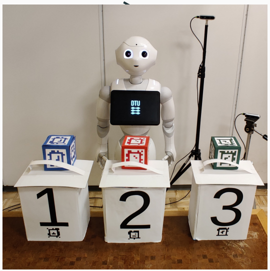

[Dissing and Bolander, 2020]

Implicit coordination: Coordinating actions without prior agreements, explicit communication or fixed protocols

Illustration by ChatGPT 5.5

Implicit coordination

Definition

A plan $i_1{:}\alpha_1, \ldots i_n{:}\alpha_n$ for an epistemic planning task $(s_0, \mathcal{A}, \varphi_g)$ is called **implicitly coordinated** if

$$s_0 \models K_{i_1} ([i_1{:}\alpha_1]) K_{i_2} ([i_2{:}\alpha_2]) \cdots K_{i_n} ([i_n{:}\alpha_n]) \varphi_g$$

Example

If the human (h) and the robot (r) shares the goal $\varphi_g = \mathit{Holding}(h, \mathit{cube})$, and the robot finds the cube first, implicit coordination guarantees that it will tell the location to the human.

Implicit coordination generalises to policies [Bolander et al., 2018].

[Engesser et al., 2017, Bolander et al., 2018]

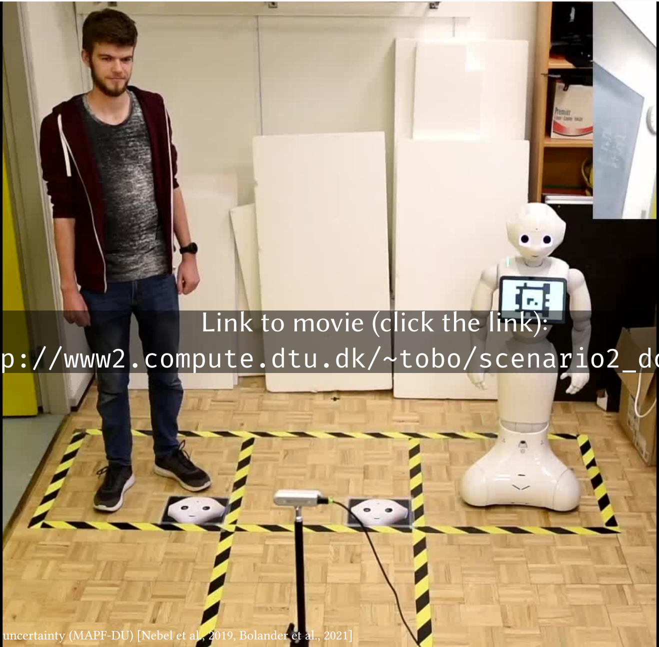

Link to movie (click the link):
http://www2.compute.dtu.dk/~tobo/scenario2_double.mp4

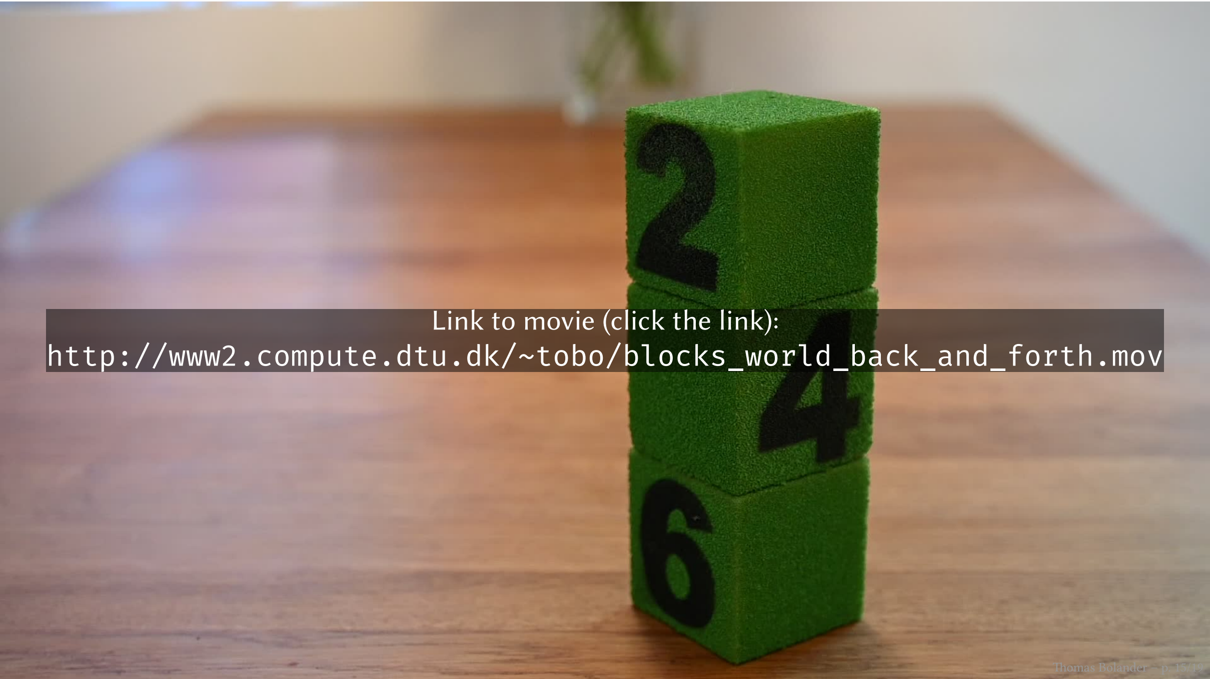

Link to movie (click the link):
http://www2.compute.dtu.dk/~tobo/blocks_world_back_and_forth.mov

There's a problem...

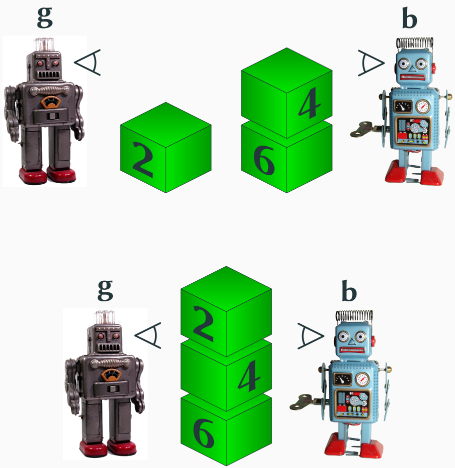

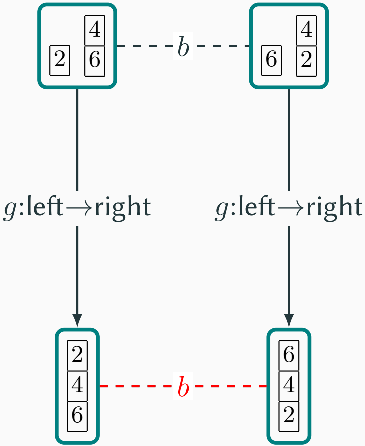

The blue agent didn't learn anything. Why not? Well, since the simple notion of implicit coordination only assumes *rationality in the future*, not *rationality in the past*.

Implicit coordination with forward induction

With the original approach to implicit coordination [Bolander et al., 2018], implicit coordination didn't guarantee successful executions! We added a notion of forward induction, inspired by epistemic game theory [Battigalli, 1997], giving us:

Theorem (Bolander and Romero, in preparation)

Implicit coordination with forward induction guarantees succesful coordination/executions. (Individual policies are coordinated: any execution of the induced non-deterministic joint policy is guaranteed to terminate in a goal state.)

Theorem (Bolander and Romero, in preparation)

Standard goal recognition by filtering is subsumed by forward induction.

$$\begin{array}{|ll|} \hline \text{pre}: & \\ \hline \text{post}: & G_a(g_1) := \top \\ \hline \end{array} \;\overset{c}{-\,-\,-}\; \begin{array}{|ll|} \hline \text{pre}: & \\ \hline \text{post}: & G_a(g_2) := \top \\ \hline \end{array} \;\cdots\cdots\; \begin{array}{|ll|} \hline \text{pre}: & \\ \hline \text{post}: & G_a(g_n) := \top \\ \hline \end{array}$$

Computational complexity of epistemic planning

In automated planning, computational complexity is most often measured by the **plan existence problem**: Given a planning task $T = (s_0, \mathcal{A}, \varphi_g)$, does it have a solution?

Theorem [Bolander et al., 2020]

1. The plan existence problem for epistemic planning tasks is decidable when pre- and post-conditions of actions are propositional.
2. It is undecidable already when either pre- or post-conditions are allowed to be epistemic formulas of modal depth 1 (and the other conditions still being propositional).

Open problem: Case of preconditions of modal depth 1 and no postconditions. Antonis Achilleos put a paper draft on Arxiv in April 2026 claiming to prove undecidability.

Taming complexity: *Depth-bounded epistemic planning* [Bolander et al., 2025] restricts agents' reasoning depth to a fixed k (by k-bisimulation contractions).
Iterative-deepening algorithm is sound and complete (and practically efficient).

Other recent and current work in epistemic planning

- **EPDDL**: A PDDL-like representation language for epistemic planning [Burigana and Fabiano, 2026].
- A new **International Epistemic Planning Competition**. First edition hosted at ICAPS 2026. `https://sites.google.com/view/epistemic-competition/`
- DEL-based epistemic planning as presented here is **model-theoretic** (states are semantic objects). **Sentential approaches** also exist (states are syntactic objects: sets of formulas). For an overview, see the AIJ special issue on epistemic planning [Belle et al., 2023].
- New research project: **Attention in Epistemic Planning (AiEP)**, 2025-2029. `https://www2.compute.dtu.dk/~tobo/aiep_webpage.html`

References i

Baltag, A., Moss, L. S., and Solecki, S. (1998).
The logic of public announcements and common knowledge and private suspicions.
In Gilboa, I., editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 43–56. Morgan Kaufmann.

Baltag, A. and Smets, S. (2008).
A qualitative theory of dynamic interactive belief revision.
In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors, *Logic and the Foundations of Game and Decision Theory (LOFT7)*, volume 3 of *Texts in Logic and Games*, pages 13–60. Amsterdam University Press.

Battigalli, P. (1997).
On rationalizability in extensive games.
Journal of Economic Theory, 74(1):40–61.

Belle, V., Bolander, T., Herzig, A., and Nebel, B. (2023).
Epistemic planning: Perspectives on the special issue.
Artificial Intelligence, 316:103842.

Bolander, T. (2018).
Seeing is believing: Formalising false-belief tasks in dynamic epistemic logic.
In *Outstanding Contributions to Logic*, number 12, pages 207–236. 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.

References ii

Bolander, T., Burigana, A., and Montali, M. (2025).
Depth-bounded epistemic planning.
In *Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning*, KR '25.

Bolander, T., Charrier, T., Pinchinat, S., and Schwarzentruber, F. (2020).
DEL-based epistemic planning: Decidability and complexity.
Artificial Intelligence, 287:1–34.

Bolander, T., Dissing, L., and Herrmann, N. (2021).
DEL-based epistemic planning for human-robot collaboration: Theory and implementation.
In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR 2021)*.

Bolander, T., Engesser, T., Mattmüller, R., and Nebel, B. (2018).
Better eager than lazy? how agent types impact the successfulness of implicit coordination.
In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018)*. AAAI Press.

Bolander, T. and Grosinger, H. J. (2026).
Epistemic proactivity via reasoning about beliefs and expectations using plausibility models.
Journal of Logic, Language and Information, pages 1–42.

Burigana, A. and Fabiano, F. (2026).
The epistemic planning domain definition language: Official guideline.
arXiv preprint arXiv:2601.20969.

Dissing, L. and Bolander, T. (2020).
Implementing Theory of Mind on a robot using Dynamic Epistemic Logic.
In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*.

Engesser, T., Bolander, T., Mattmüller, R., and Nebel, B. (2017).
Cooperative epistemic multi-agent planning for implicit coordination.
In *Proceedings of Methods for Modalities*, number 243 in Electronic Proceedings in Theoretical Computer Science, pages 75–90.

Ghallab, M., Nau, D. S., and Traverso, P. (2004).
Automated Planning: Theory and Practice.
Morgan Kaufmann.

Nebel, B., Bolander, T., Engesser, T., and Mattmüller, R. (2019).
Implicitly coordinated multi-agent path finding under destination uncertainty: Success guarantees and computational complexity.
Journal of Artificial Intelligence Research, 64:497–527.

Pieper, J., Engesser, T., and Nebel, B. (2025).
Towards implicit coordination planning with knowledge and belief.
In *Festschrift for Andreas Herzig on the Occasion of his 65th Birthday: Essays in Honor of Andi.* College Publications.

van Ditmarsch, H. and Kooi, B. (2008).
Semantic results for ontic and epistemic change.
In Bonanno, G., van der Hoek, W., and Wooldridge, M., editors, *Logic and the Foundation of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games 3, pages 87–117. Amsterdam University Press.

Warneken, F. and Tomasello, M. (2006).
Cabinet task.
Max Planck Institute for Evolutionary Anthropology.
`https://www.eva.mpg.de/fileadmin/content_files/psychology/videos/children_cabinet.mpg` [Accessed: 24 November 2021].