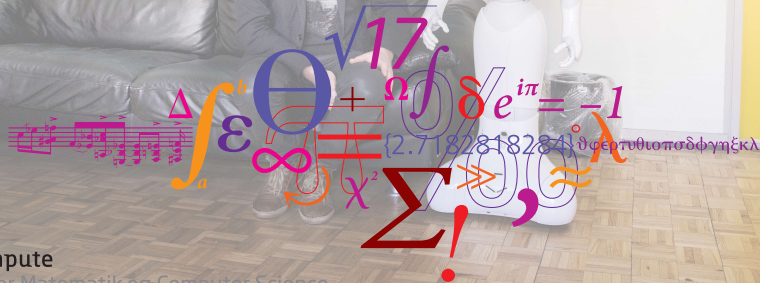


From Dynamic Epistemic Logic to Socially Intelligent Robots

Thomas Bolander, DTU Compute

WoLLIC 2023



Automated planning: Classical planning tasks

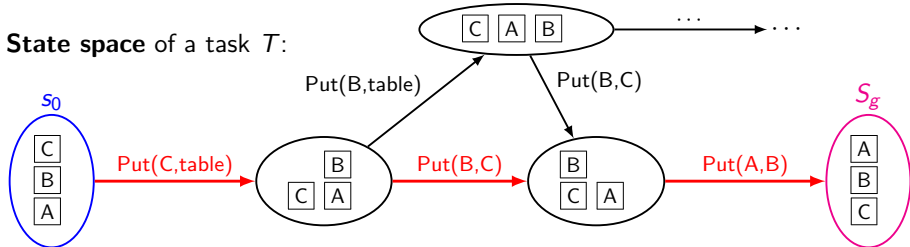
Definition (Adopted from [Ghallab et al., 2004])

A **classical planning task** (over states S) is $T = (s_0, A, \circ, S_g)$, where

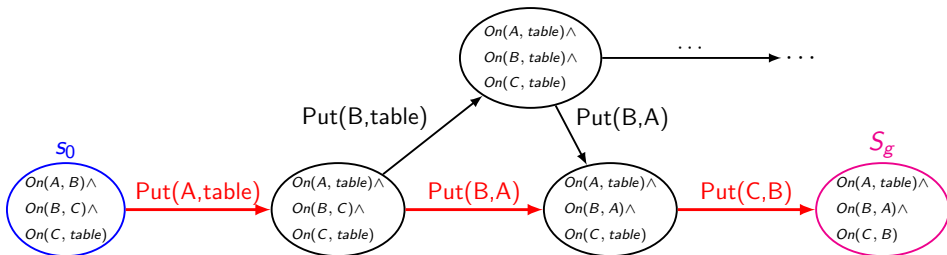
- $s_0 \in S$ is an **initial state**.
- A is a finite set of available **actions**.
- \circ is a state-transition operator: for $s \in S, \alpha \in A$, either $s \circ \alpha \in S$ or $s \circ \alpha$ is undefined (and in that case we say α is **inapplicable** in s).
- $S_g \subseteq S$ is a set of **goal states**.

A **solution** to a classical planning task (s_0, A, \circ, S_g) is a sequence of actions (a **plan**) $\pi = \alpha_1, \dots, \alpha_n$ from A such that $s_0 \circ \alpha_1 \circ \dots \circ \alpha_n \in S_g$.

State space of a task T :



Action schemas

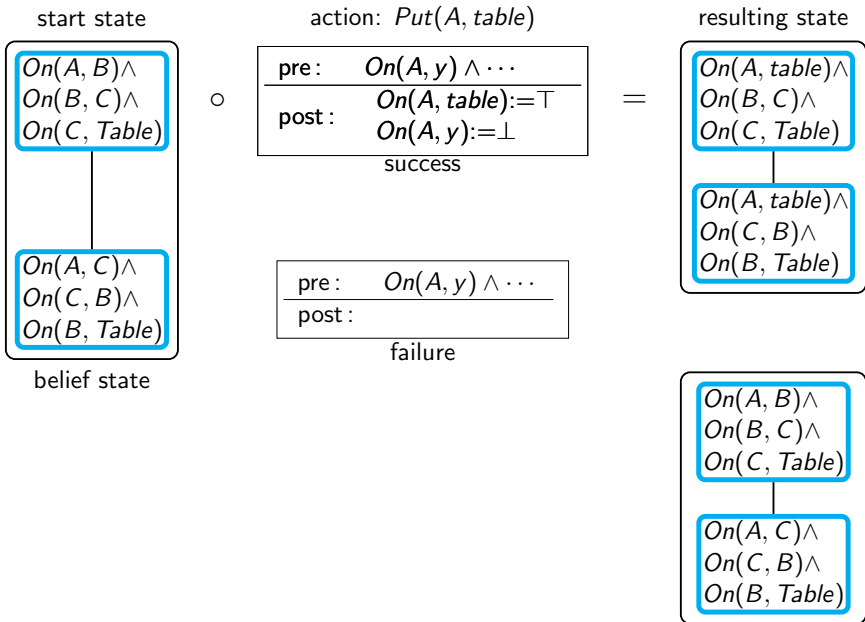


Action schema describing the $Put(x, y)$ action for put object x on top of object y :

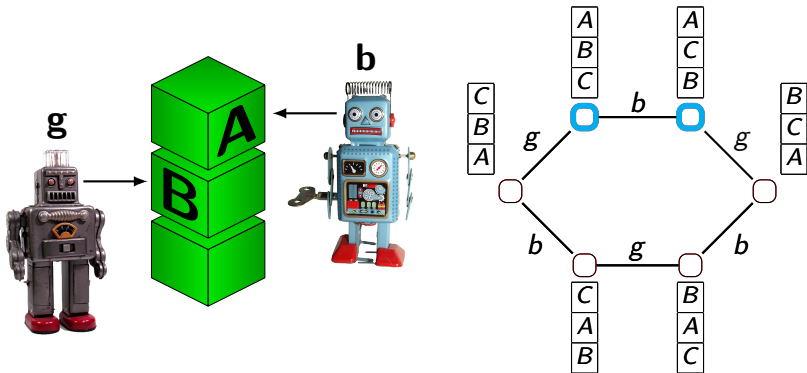
ACTION : $Put(x, y)$
PRECONDITION : $On(x, z) \wedge \dots$
EFFECT : $On(x, y) \wedge \neg On(x, z)$

pre :	$On(x, z) \wedge \dots$
post :	$On(x, y) := \top$ $On(x, z) := \perp$

Adding non-determinism and partial observability



Multiagent case: States as 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 planning agent).

Agent *b*: “Which letter does the middle block have?”
(Public Announcement Logic, PAL [Plaza, 1989])

Implicit coordination

Suppose the agents have a **joint goal** (like stacking blocks in reverse alphabetic order), and **plan independently**. We use notation $i:\alpha$ for “agent i does α ”.

An fully observant agent c might form this plan (where the blocks are numbered 1–3 according to their initial stacking order):

$b:Put(1, table), g:Put(2, 1), b:Put(3, 2)$.

However, it's not a **verifiable** solution by b and g .

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 .

Redefined solution concept: A **plan** is $i_1:\alpha_1, \dots, i_n:\alpha_n$ such that

$$(\dots((s_0^{i_1} \circ i_1:\alpha_1)^{i_2} \circ i_2:\alpha_2)^{i_3} \circ \dots \circ i_{n-1}:\alpha_{n-1})^{i_n} \circ i_n:\alpha_n \in S_g$$

Problem: Only assumes other agents to be rational in the future.

Solution: Introduce forward induction (work in progress).

[Engesser et al., 2017, Bolander et al., 2016]

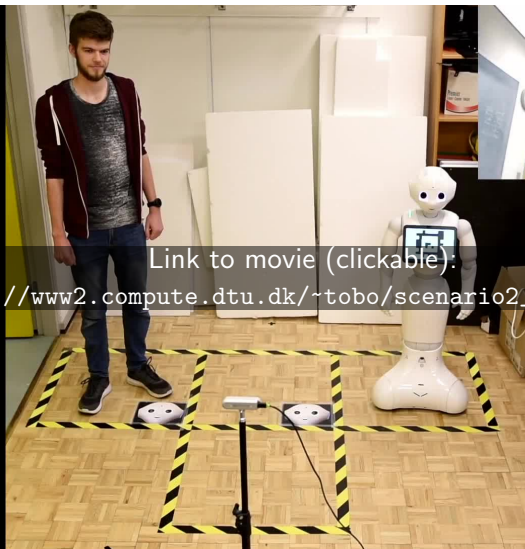
F	A	A	0									A
F									4	F	F	A
G	B	B	1									D
G									5	G	G	D
H	C	C	2									C
H									6	H	H	C
I	D	D	3									B
I									7	I	I	B

Link to movie (clickable):

http://www2.compute.dtu.dk/~tobo/MAREgExAZ_02285_level_vert.mov

F	A	A	0									A
F									4	F	F	A
G	B	B	1									D
G									5	G	G	D
H	C	C	2									C
H									6	H	H	C
I	D	D	3									B
I									7	I	I	B

Multi-agent pathfinding with destination uncertainty

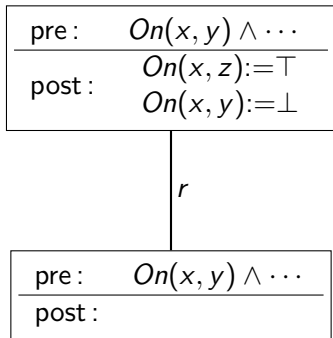


Link to movie (clickable):

http://www2.compute.dtu.dk/~tobo/scenario2_double.mp4

[Nebel et al., 2019, Bolander et al., 2021]

Introducing partially observable actions



This is an **event model** of dynamic epistemic logic (DEL)...

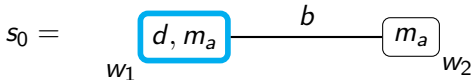
Dynamic epistemic logic (DEL) via example: The coordinated attack problem

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**:

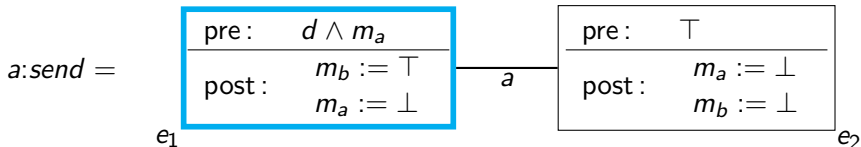


Nodes are **worlds**, edges are **indistinguishability edges** (reflexive loops not shown).

Event models of 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**):

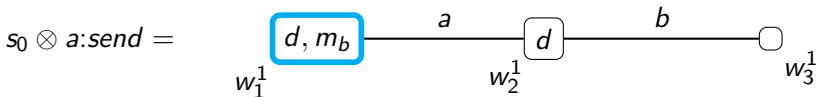
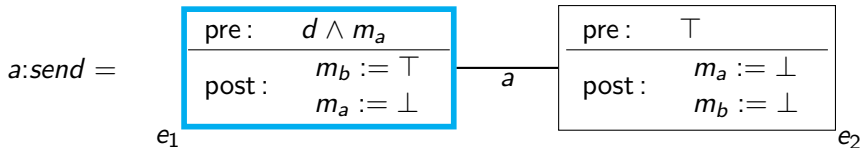
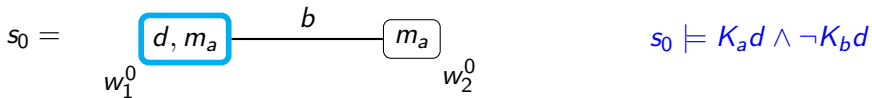


And symmetrically an epistemic action $b:send$. We read $i:\alpha$ as “agent i does α ”.

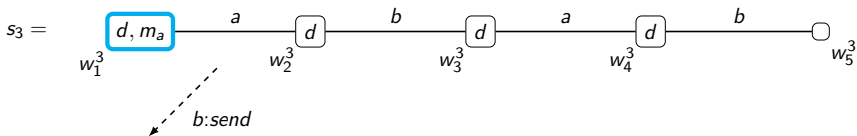
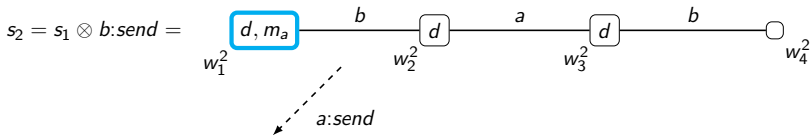
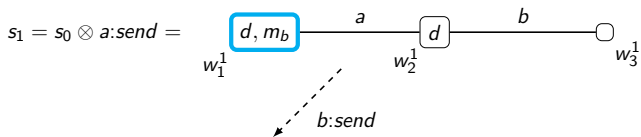
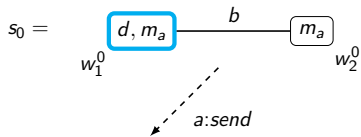
Nodes are **events**, and each event has a **precondition** and a **postcondition** (effect). The precondition is an epistemic formula and the postcondition is a conjunction of literals.

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008, Bolander et al., 2021]

The product update in dynamic epistemic logic



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



Epistemic planning tasks

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

Definition. A (sequential) **solution** to a planning task $T = (s_0, A, \varphi_g)$ 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 $s_0 \otimes \alpha_1 \otimes \dots \otimes \alpha_{i-1}$ and

$$s_0 \otimes \alpha_1 \otimes \alpha_2 \otimes \dots \otimes \alpha_n \models \varphi_g.$$

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:

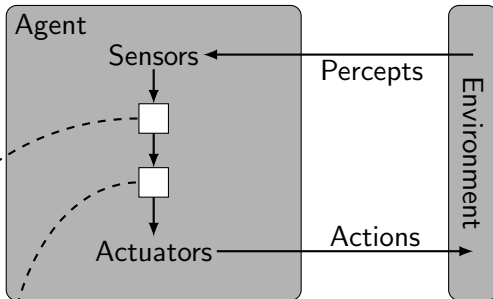
1. $T = (s_0, A, Cd)$, where C denotes common knowledge. It has no solution.
2. $T = (s_0, A, E^n d)$, where E denotes “everybody knows” and $n \geq 1$. It has a solution of length n .

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/komdigital_pepper_video.mov

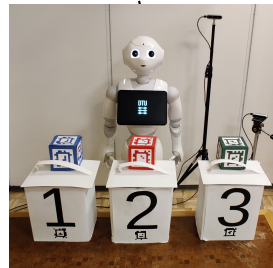
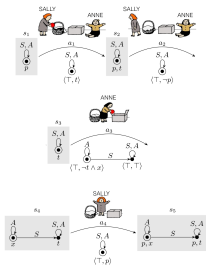
Thomas Bolander, Professor
DTU Compute
Technical University of Denmark

KomDigital: R2DTU – A Pepper robot, 25 November 2020 [?]

perception layer:
 deep neural
 networks
 (subsymbolic)



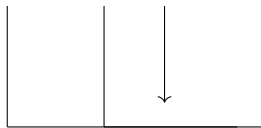
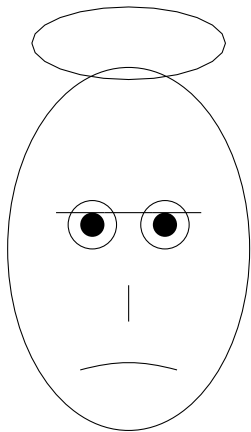
cognition layer:
 DEL + planning
 (symbolic)



Why not just make ChatGPT or GPT-4 do it all?

GPT-4 can do amazing things.

Example. My Input to GPT-4: “Provide tikz code for an illustration of Kurt Gödel, the famous logician who proved the incompleteness theorems.”



Kurt Gödel

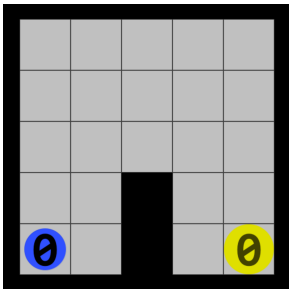
TO

Suppose you're controlling a robot in a grid world. The robot can in each move move one cell north, east, west or south. The grid cells are labelled by coordinates with (1,1) in the lower left corner. Assume you want to move the robot from (1,1) to (5,1) and assume there are walls in cell (3,1) and (3,2). The robot can't move through wall cells. Please devise a plan for the robot, i.e., a sequence of moves that brings it from the initial state, (1,1), to the goal state, (5,1).



One possible solution could be:

1. Move the robot east to (2,1).
2. Move the robot east to (3,1).
3. Move the robot north to (3,2).
4. Move the robot west to (2,2).
5. Move the robot south to (2,1).
6. Move the robot east to (3,1).
7. Repeat steps 4 to 6 until the robot reaches (5,1).



Plan existence problem

Definition. $\mathcal{T}(m, n)$ is the class of epistemic planning tasks where all actions have preconditions of modal depth $\leq m$ and postconditions of modal depth $\leq n$. We use $n = -1$ to denote the case without postconditions.

in $\mathcal{T}(0, 0)$

pre:	$p \wedge q$
post:	$p := \perp$ $r := \top$

in $\mathcal{T}(2, -1)$

pre:	$K_b K_g On(C, B)$
post:	

in $\mathcal{T}(1, 1)$

pre:	$K_1 p$
post:	$p := K_2 q$

Definition. $\text{PlanEx-}\mathcal{T}(m, n)$ is the following decision problem: Given a planning task $T \in \mathcal{T}(m, n)$, does T have a solution?

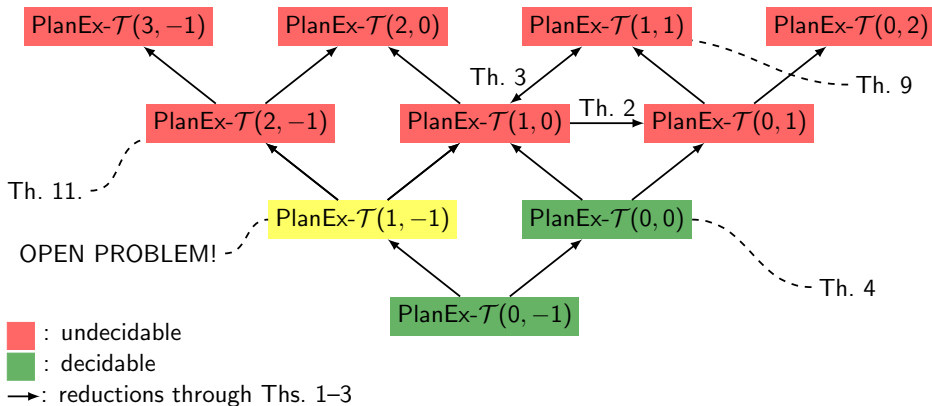
[Bolander et al., 2020]

The border between decidability and undecidability

Theorem 1. $\text{PlanEx-}\mathcal{T}(m, n) \leq^P \text{PlanEx-}\mathcal{T}(m+k, n+l)$.

Theorem 2. $\text{PlanEx-}\mathcal{T}(m, n) \leq^P \text{PlanEx-}\mathcal{T}(0, 1)$.

Theorem 3. $\text{PlanEx-}\mathcal{T}(m, n) \leq^P \text{PlanEx-}\mathcal{T}(1, 0)$.



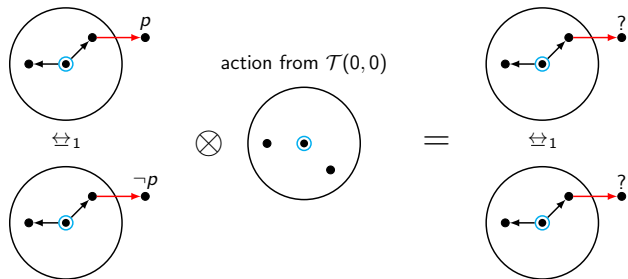
[Bolander and Andersen, 2011, Aucher and Bolander, 2013, Yu et al., 2013, Charrier et al., 2016, Cong et al., 2018, Bolander et al., 2020]

Decidability theorem

k -bisimilarity: Satisfying back and forth conditions of bisimilarity up to depth k . Guarantees modal equivalence up to modal depth k .

Theorem 4. *PlanEx- $\mathcal{T}(0,0)$ is decidable.*

Proof idea: k -bisimilarity is preserved when doing product update with epistemic actions having propositional pre- and post-conditions [Yu et al., 2013]; intuitively because the events of such actions cannot look deeper into the model.



$\mathcal{T}(0,0)$ examples

public announcements
of propositional facts:

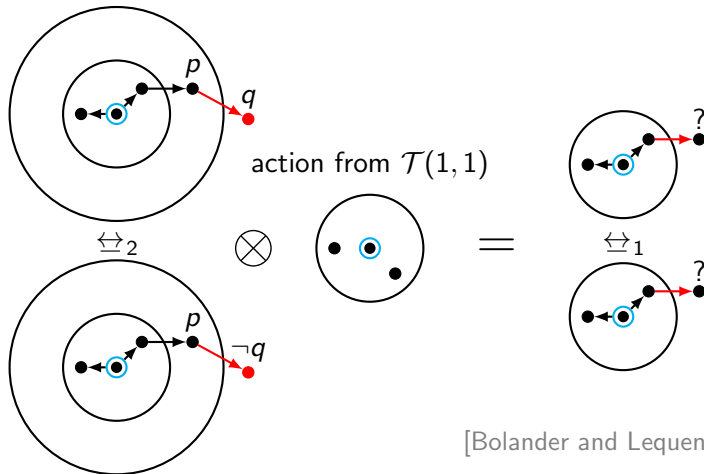
pre :	$p \wedge q$
post :	

STRIPS actions:

pre :	$p \wedge q$
post :	$p := \perp$ $r := \top$

Generalising the k -bisimilarity preservation result

Proposition 1. Suppose s and s' are k -bisimilar and α is an action of $\mathcal{T}(m, n)$. Then $s \otimes \alpha$ and $s' \otimes \alpha$ are $(k - \max\{m, n\})$ -bisimilar.



[Bolander and Lequen, 2022]

Depth-bounded epistemic planning (w. in progress)

Planning algorithm $\text{SEARCH}(T, k)$ with depth-bound k : breadth-first search (BFS) through the state space, exploiting Proposition 1:

- Whenever we apply action α to state s , we afterward do the k -bisimulation contraction, where k is the maximal bound guaranteeing preservation of k -bisimilarity.
- We terminate any path satisfying $k \leq \text{modal-depth}(\varphi_g)$.

Parameters of planning task T (we study parameterised complexity).

a: number of agents.

c: maximal modal depth of preconditions of actions.

o: modal depth of goal formula.

p: number of propositional variables.

u: maximal length of plan.

Soundness. If $\text{SEARCH}(T, k)$ returns π , then π is a solution to T .

Completeness. If T has a solution, it will be found by $\text{SEARCH}(T, k)$ whenever $k \geq cu + o$.

Complexity. $\text{SEARCH}(T, k)$ runs in time $\exp_2^{cu+o+1} \max\{a, p\}$.

Appendix: References I



Aucher, G. and Bolander, T. (2013).

Undecidability in Epistemic Planning.

In Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI) pp. 27–33,.



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.



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 Schwarzentruher, 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. (2016).

Better Eager Than Lazy? How Agent Types Impact the Successfulness of Implicit Coordination.

In Distributed and Multi-Agent Planning (DMAP-16) pp. 42–49,.



Bolander, T. and Lequen, A. (2022).

Parameterized complexity of dynamic belief updates: A complete map.

Journal of Logic and Computation 0, 0–0.

Appendix: References II



Charrier, T., Maubert, B. and Schwarzenrüber, 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, NY, USA, 9-15 July 2016 pp. 1030–1036,.



Cong, S. L., Pinchinat, S. and Schwarzenrüber, F. (2018).

Small Undecidable Problems in Epistemic Planning.

In Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden. pp. 4780–4786,.



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 pp. 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.



Plaza, J. (1989).

Logics of public announcements.

In Proceedings 4th International Symposium on Methodologies for Intelligent Systems pp. 201–216,.

Appendix: References III



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.



Yu, Q., Wen, X. and Liu, Y. (2013).

Multi-agent epistemic explanatory diagnosis via reasoning about actions.

In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)* pp. 27–33,.