From Dynamic Epistemic Logic to Socially Intelligent Robots

Thomas Bolander, DTU Compute

LORI, 26 October 2023

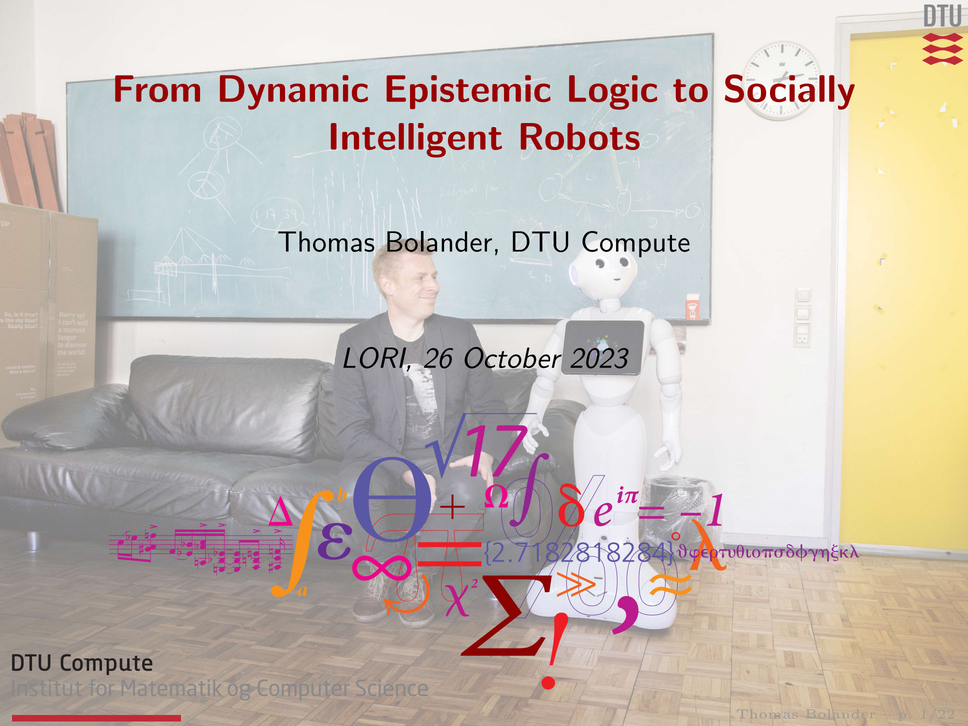

DTU Compute
Institut for Matematik og Computer Science

Epistemic planning =

automated *planning* (AI) + *epistemic* reasoning (epistemic logic)

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**.

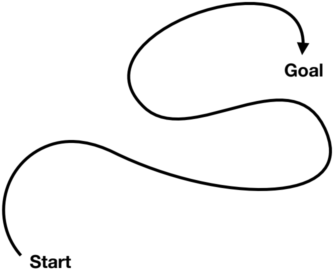

Automated planning

$+$

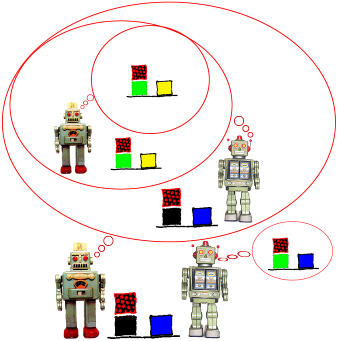

Epistemic reasoning about the mental states of others

Classical automated planning: single agent, full observability

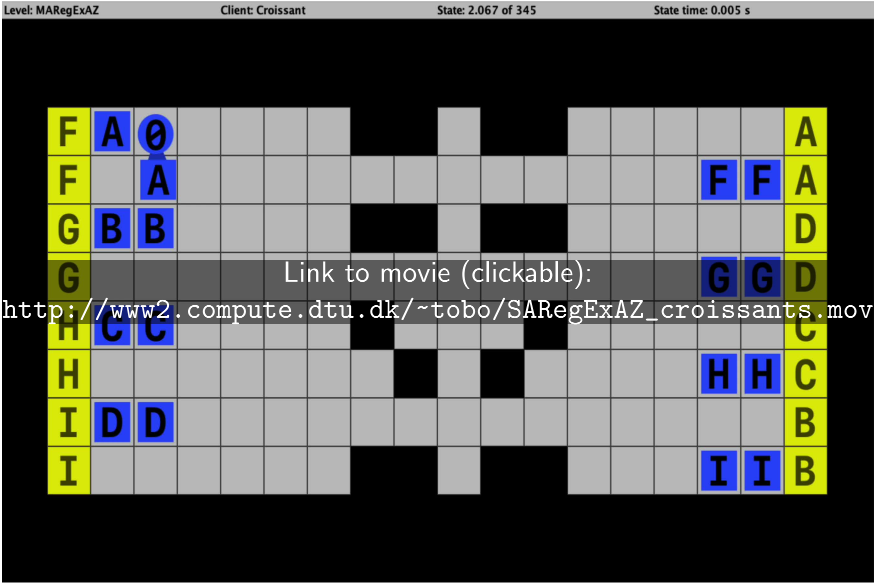

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/SARegExAZ_croissants.mov

Classical automated planning: state space search and domain descriptions

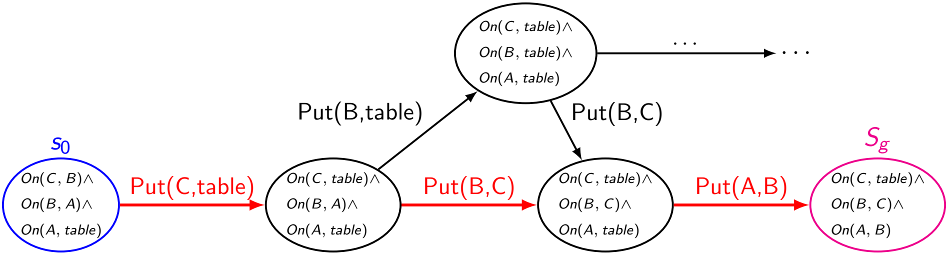

Action schema describing the $Put(x,y)$ action for put object x on top of object y:

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

pre :	$On(x,z) \wedge \cdots$
post :	$On(x,y) := \top$ $On(x,z) := \bot$

[Ghallab et al., 2004, Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

Amazon warehouse robots

Link to movie:
http://www2.compute.dtu.dk/~tobo/amazon_kiva.mov

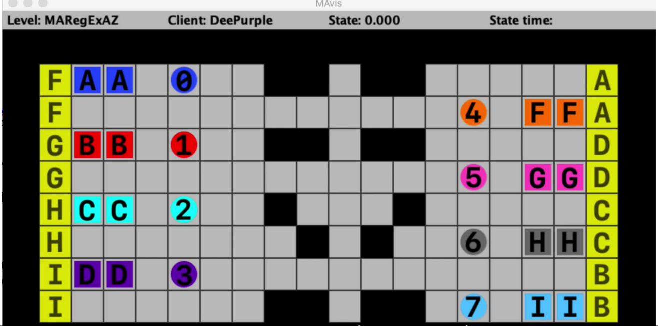

Link to movie (clickable):

http://www2.compute.dtu.dk/~tobo/MARegExAZ_02285_level_vert.mov

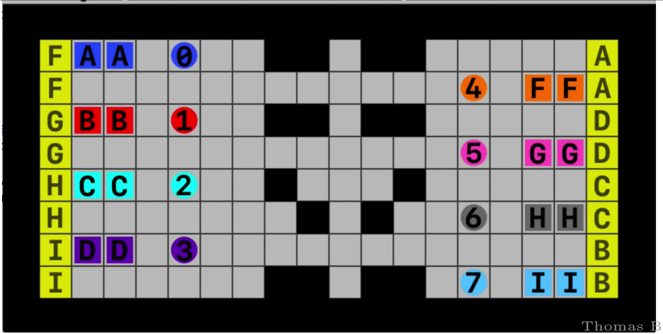

Multiagent case: States as S5 Kripke models

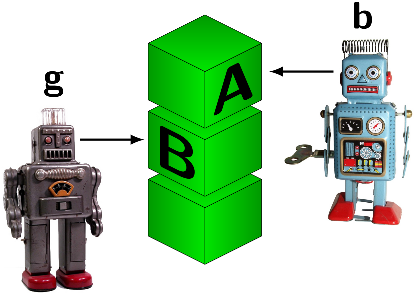

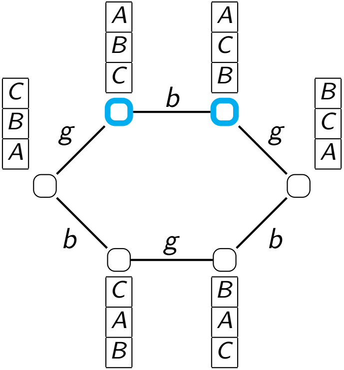

Epistemic states: Multi-pointed epistemic models of multi-agent S5. Nodes are **worlds**, edges are **indistinguishability relations**.
Designated worlds: (those considered possible by planning agent).

Agent *b*: "Which letter does the middle block have?"

Dynamic epistemic logic (DEL) by example: product update

$s_0 =$

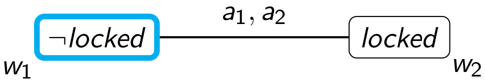

$a_1{:}open =$

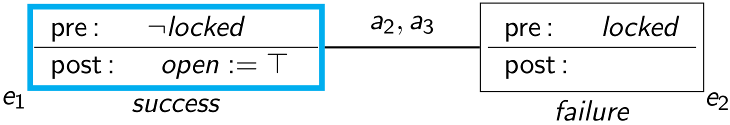

$s_0 \otimes a_1{:}open =$

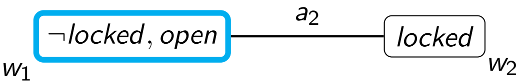

[Baltag et al., 1998, van Ditmarsch and Kooi, 2008]

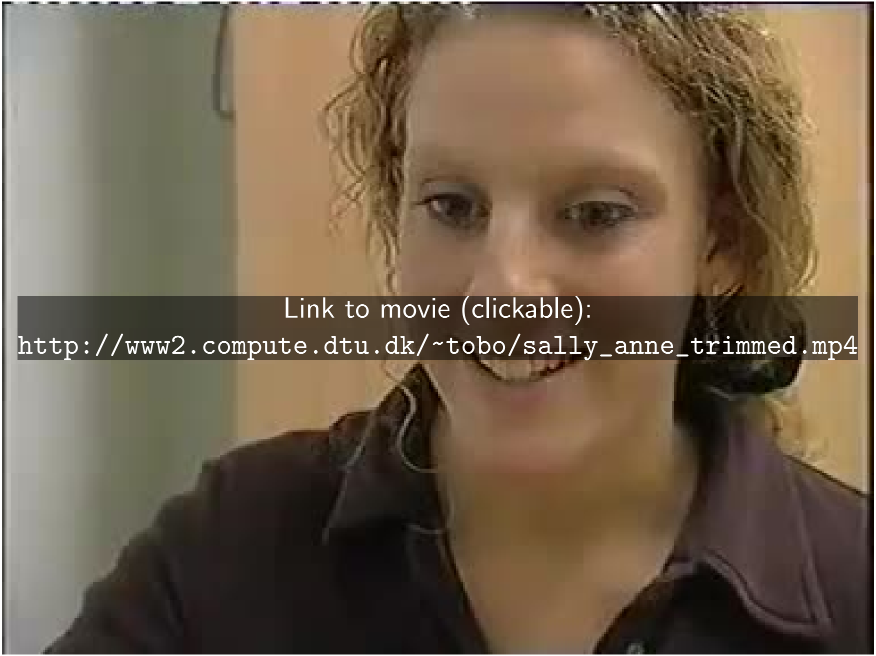

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/sally_anne_trimmed.mp4

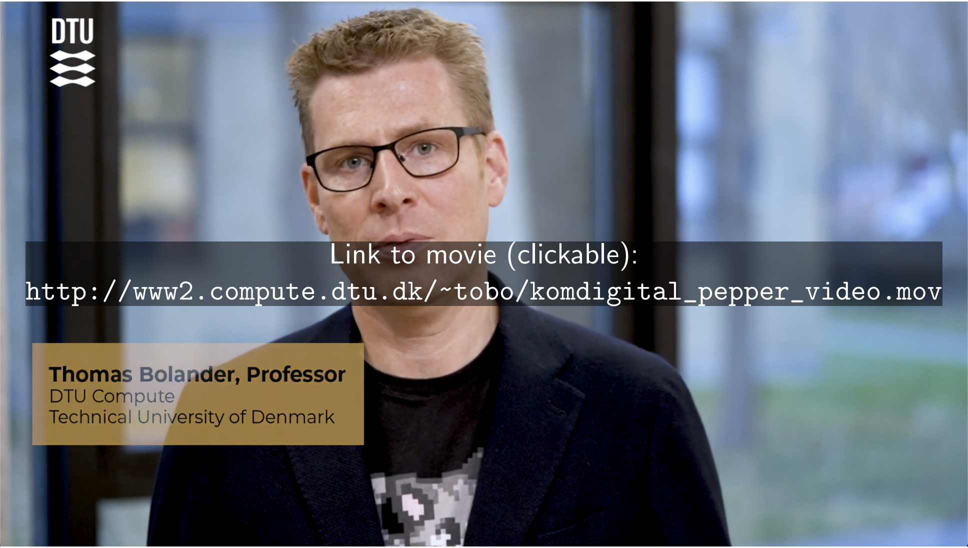

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/komdigital_pepper_video.mov

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

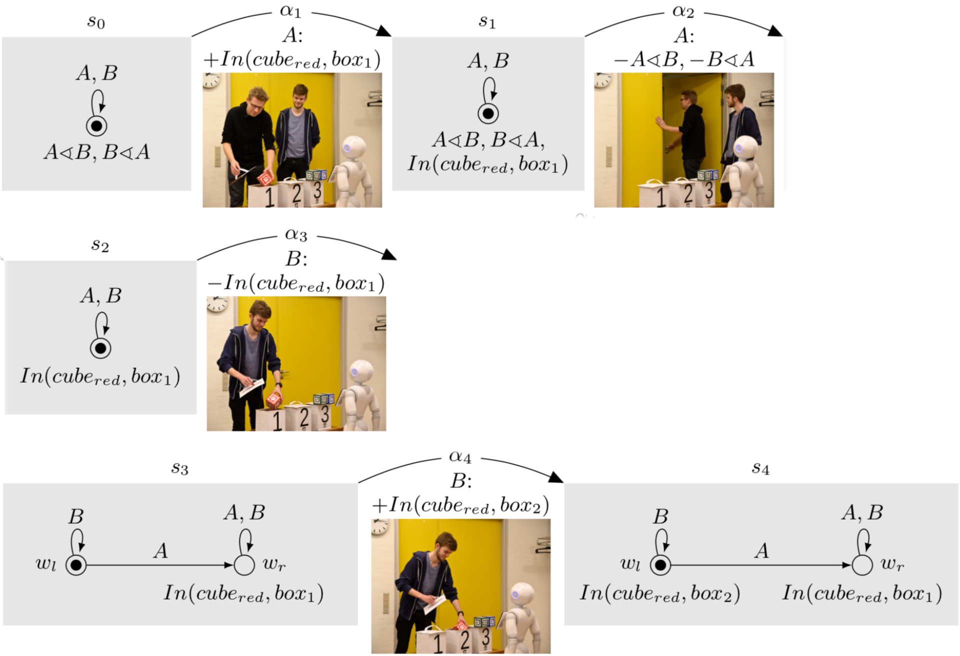
s_0
A, B
$A \lhd B, B \lhd A$
α_1
A:
$+In(cube_{red}, box_1)$
s_1
A, B
$A \lhd B, B \lhd A,$
$In(cube_{red}, box_1)$
α_2
A:
$-A \lhd B, -B \lhd A$
s_2
A, B
$In(cube_{red}, box_1)$
α_3
B:
$-In(cube_{red}, box_1)$
s_3
B
A, B
w_l
A
w_r
$In(cube_{red}, box_1)$
α_4
B:
$+In(cube_{red}, box_2)$
s_4
B
A, B
w_l
A
w_r
$In(cube_{red}, box_2)$
$In(cube_{red}, box_1)$

Planning based on DEL: epistemic planning tasks

Definition. An **(epistemic) planning task** $T = (s_0, A, \varphi_g)$ consists of

- A multipointed Kripke model s_0 called the **initial state**.
- A finite set of multipointed event models A called **actions**.
- A **goal formula** φ_g of epistemic logic.

Definition. A (sequential) **solution** to a planning task $T = (s_0, A, \varphi_g)$ is a sequence of actions $\alpha_1, \alpha_2, \ldots, \alpha_n$ 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

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

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

$$s_0 \models (\!(\alpha_1)\!)(\!(\alpha_2)\!)\cdots(\!(\alpha_n)\!)\varphi_g.$$

Definition. A solution $i_1{:}\alpha_1, \ldots i_n{:}\alpha_n$ (using notation $i{:}\alpha$ for agent i performing action α) is **implicitly coordinated** if it furthermore holds that :

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

[Bolander and Andersen, 2011, Engesser et al., 2017, Bolander et al., 2020]

Implicit coordination: multi-agent pathfinding with destination uncertainty

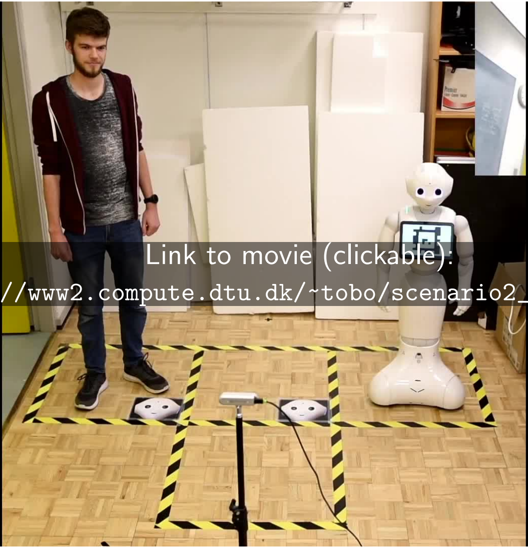

Link to movie (clickable):
http://www2.compute.dtu.dk/~tobo/scenario2_double.mp4

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

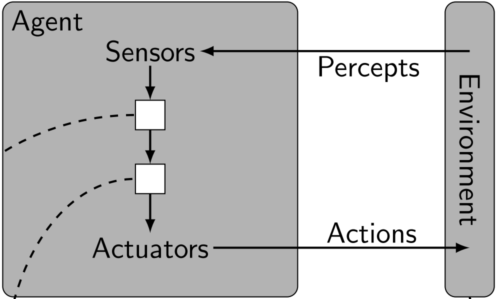

perception layer:
deep neural networks (subsymbolic)

cognition layer:
DEL + planning (symbolic)

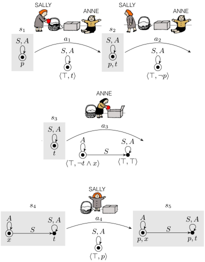

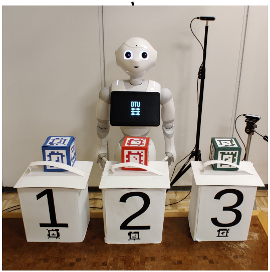

But...

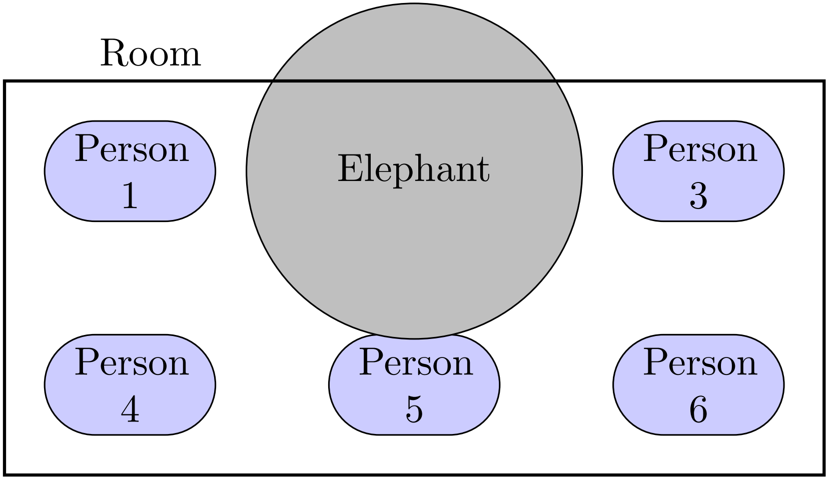

My prompt to GPT-4: “Please make a tikz illustration of the elephant in the room.”

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).

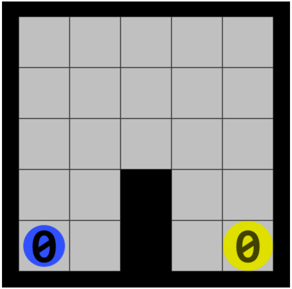

On the Planning Abilities of Large Language Models (A Critical Investigation with a Proposed Benchmark)

Karthik Valmeekam
School of Computing & AI
Arizona State University, Tempe.
kvalmeek@asu.edu

Sarath Sreedharan*
Department of Computer Science,
Colorado State University, Fort Collins.
sarath.sreedharan@colostate.edu

Alberto Olmo
School of Computing & AI
Arizona State University, Tempe.
aolmoher@asu.edu

Subba
School
Arizona St
r

Abstract

Intrigued by the claims of emergent reasoning capabilities in LLMs t
in this paper, we set out to investigate their planning capabilities. We
LLMs are by themselves in generating and validating simple plans i
(of the type that humans are generally quite good at) and (2) how good
heuristic guidance for other agents–either AI planners or human plann
investigate these questions in a systematic rather than anecdotal ma
benchmark suite based on the kinds of domains employed in the Inter
On this benchmark, we evaluate LLMs in three modes: *autonomous, h*
Our results show that LLM's ability to autonomously generate exe
averaging only about 3% success rate. The heuristic and human-in-
more promise. In addition to these results, we also make our benchmar
to support investigations by research community.

February 2023

Clever Hans or Neural Theory of Mind? Stress Testing Social Reasoning in Large Language Models

Natalie Shapira[1] **Mosh Levy***[1] **Seyed Hossein Alavi***[2,3] **Xuhui Zhou***[4]
Yejin Choi[5,6] **Yoav Goldberg**[1,5] **Maarten Sap**[4,5] **Vered Shwartz**[2,3]
[1] Bar-Ilan University [2] University of British Columbia
[3] Vector Institute for AI [4] Carnegie Mellon University
[5] Allen Institute for Artificial Intelligence [6] University of Washington
nd1234@gmail.com

Abstract

The escalating debate on AI's capabilities warrants developing reliable metrics to assess machine "intelligence." Recently, many anecdotal examples were used to suggest that newer large language models (LLMs) like ChatGPT and GPT-4 exhibit Neural Theory-of-Mind (N-ToM); however, prior work reached conflicting conclusions regarding those abilities. We investigate the extent of LLMs' N-ToM through an extensive evaluation on 6 tasks and find that while LLMs exhibit certain N-ToM abilities, this behavior is far from being robust. We further examine the factors impacting performance on N-ToM tasks and discover that LLMs struggle with adversarial examples, indicating reliance on shallow heuristics rather than robust ToM abilities. We caution against drawing conclusions from anecdotal examples, limited benchmark testing, and using human-designed psychological tests to evaluate models.

Two recent papers addressed whether Large Language Models (LLMs; Brown et al., 2020; Bommasani et al., 2021; Zhao et al., 2023) have a ToM, and came to opposite conclusions: Sap et al. (2022) shows they lack this ability and Kosinski (2023) claims this ability has emerged in the newer models spontaneously. The latter was criticized for its flawed methodology (Marcus and Davis, 2023). Ullman (2023) further showed that simple changes to the ToM questions break LLMs. But to paraphrase the saying, hype gets halfway around the world before rigorous experiments put on their boots; other researchers continue to spread the word about N-ToM, claiming that GPT-4 *"has a very advanced level of theory of mind"* based on a few anecdotal examples (Bubeck et al., 2023).

Do LLMs have robust N-ToM? This paper aims to address the discrepancy and limited scope of previous work (that each tested 2 tasks) by performing

May 2023

Plan existence problem

Definition. $\mathcal{T}(m, n)$ is the class of epistemic planning tasks where all actions have

- Preconditions of modal depth $\leq m$
- Postconditions of modal depth $\leq n$.

We use $n = -1$ to denote the case without postconditions.

Examples.

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

pre :	$p \wedge q$
post :	$p{:=}\bot$ $r{:=}\top$

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

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

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

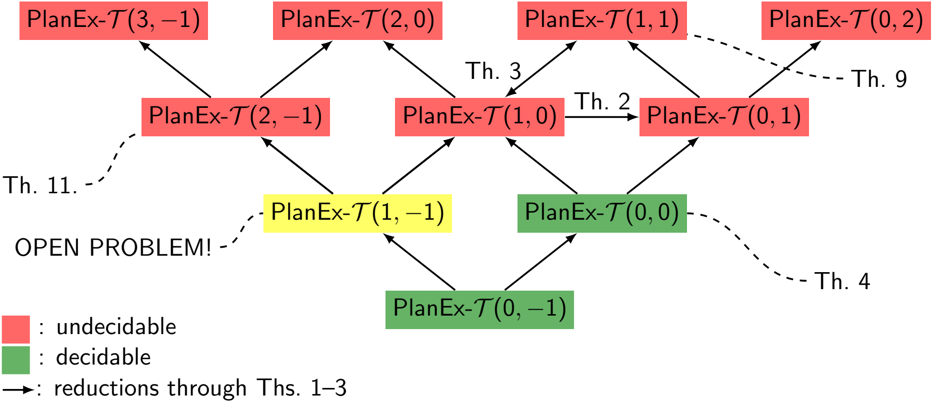

[Bolander et al., 2020]

k-**bisimilarity**, $\underline{\leftrightarrow}_k$: Models satisfying back and forth conditions of bisimilarity up to depth k. Gives modal equivalence to modal depth k.

Theorem. *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.*

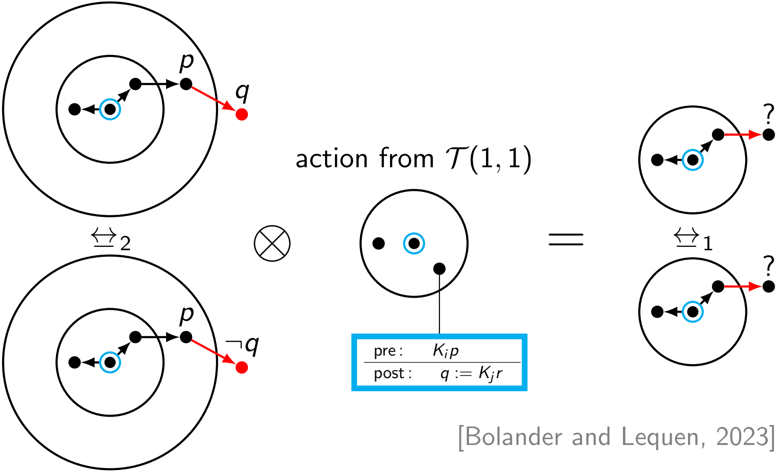

[Bolander and Lequen, 2023]

Theorem. *PlanEx-$\mathcal{T}(0,0)$ is decidable. (Orig. proof [Yu et al., 2013])*

Depth-bounded epistemic planning (w. in progress)

Planning algorithm $\textsc{Search}(T, k)$ with depth-bound k:

- Take k-bisimulation contraction of initial state s_0.
- After each product update, do l-bisimulation contraction for largest possible l (by partition refinement, using a new approach).
- If $l < \mathit{modal\text{-}depth}(\varphi_g)$, terminate the current search path.

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

a: # agents
p: # propositional variables
o: modal depth of goal formula
u: maximal length of plan
c: max. modal depth of action preconditions

Soundness and completeness. If $\textsc{Search}(T, k)$ returns π, then π is a solution to T. If T has a solution, it will be found by $\textsc{Search}(T, k)$ whenever $k \geq \mathrm{cu} + \mathrm{o}$.

Complexity. $\textsc{Search}(T, k)$ runs in time $\exp_2^{k+1} \max\{\mathrm{a}, \mathrm{p}\}$.

For any proper subset of the paramaters acopu, even plan verification is fixed-parameter intractable. [Bolander and Lequen, 2023]

Some open questions in epistemic planning

- (Un)decidability of PlanEx-$\mathcal{T}(1, -1)$.
- Other *natural* restrictions in epistemic planning, e.g. structure on formulas (suggested by Johan van Benthem).
- Theory of Mind with other notions than belief, knowledge and observability: attention, goals/intentions, etc.
- Implicit coordination done right (with forward induction, allowing for goal recognition).
- Heuristics in epistemic planning.

Appendix: 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.

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.

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. and Lequen, A. (2023).
Parameterized complexity of dynamic belief updates: A complete map.
Journal of Logic and Computation *33*, 1270–1300.

DTUdk (2020).
KomDigital: R2DTU – A Pepper robot.
Youtube.
https://www.youtube.com/watch?v=9qYCFew58uM [Accessed: 24 November 2021].

Appendix: References II

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.

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,.