

Complexity Results in Epistemic Planning

Thomas Bolander
 DTU Compute
 Tech. University of Denmark
 Copenhagen, Denmark
 tobo@dtu.dk

Martin Holm Jensen
 DTU Compute
 Tech. University of Denmark
 Copenhagen, Denmark
 martin.holm.j@gmail.com

Francois Schwarzentruber
 IRISA
 ENS Rennes
 Rennes, France
 francois.schwarzentruber@ens-rennes.fr

Abstract

Epistemic planning is a very expressive framework that extends automated planning by the incorporation of dynamic epistemic logic (DEL). We provide complexity results on the plan existence problem for multi-agent planning tasks, focusing on purely epistemic actions with propositional preconditions. We show that moving from epistemic preconditions to propositional preconditions makes it decidable, more precisely in EXPSPACE. The plan existence problem is PSPACE-complete when the underlying graphs are trees and NP-complete when they are chains (including singletons). We also show PSPACE-hardness of the plan verification problem, which strengthens previous results on the complexity of DEL model checking.

1 Introduction

An all-pervading focus of artificial intelligence (AI) is the development of rational, autonomous agents. An important trait of such an agent is that it is able to exhibit goal-directed behaviour, and this overarching aim is what is studied within the field of automated planning. At the same time, such goal-directed behaviour will naturally be confined to whatever model of the underlying domain is used. In automated planning the domain models employed are formulated using propositional logic, but in more complex settings (e.g. multi-agent domains) such models come up short due to the limited expressive power of propositional logic. By extending (or replacing) this foundational building block of automated planning we obtain a more expressive formalism for studying and developing goal-directed agents, enabling for instance an agent to reason about other agents.

For the above reasons automated planning has recently seen an influx of formalisms that are colloquially referred to as epistemic planning [Bolander and Andersen, 2011; Löwe *et al.*, 2011; Aucher and Bolander, 2013; Yu *et al.*, 2013; Andersen *et al.*, 2012]. Common to these approaches is that they take dynamic epistemic logic (DEL) [Baltag *et al.*, 1998] as the basic building block of automated planning, which greatly surpasses propositional logic in terms of expressive power. Briefly put, DEL is a modal logic with which we

can reason about the dynamics of knowledge. In the single-agent case, epistemic planning can capture non-deterministic and partially observable domains [Andersen *et al.*, 2012]. An even more interesting feature of DEL is the inherent ability to reason about multi-agent scenarios, lending itself perfectly to natural descriptions of multi-agent planning tasks.

In [Bolander and Andersen, 2011] it is shown that the plan existence problem (i.e. deciding whether a plan exists for a multi-agent planning task) is undecidable, and this remains so even when factual change is not allowed, that is, when we only allow actions that changes beliefs, not ontic facts [Aucher and Bolander, 2013]. Allowing for factual change, a decidable fragment is obtained by restricting epistemic actions to only have propositional preconditions [Yu *et al.*, 2013] (in the full framework, preconditions of actions can be arbitrary *epistemic* formulas). The computational complexity of this fragment belongs to $(d + 1)$ -EXPTIME for a goal whose modal depth is d [Maubert, 2014].

In this work we consider exclusively the plan existence problem for classes of planning tasks where preconditions are propositional (as in most automated planning formalisms) and actions are non-factual (changing only beliefs). We show this problem to be in EXPSPACE in the general case, but also identify fragments with tight complexity results. We do so by using the notion of epistemic action stabilisation [van Benthem, 2003; Miller and Moss, 2005; Sadzik, 2006], which allows us to put an upper bound on the number of times an action needs to be executed in a plan. This number depends crucially on the structural properties of the graph underlying the epistemic action. To achieve our upper bound complexity results we generalise a result of [Sadzik, 2006] on action stabilisation. We also tackle lower bounds, thereby showing a clear computational separation between these fragments.

Our contributions to the complexity of the plan existence problem are summarised in Table 1 (second column from the left), where we've also listed related contributions. The fragments we study have both a conceptual and technical motivation. Singleton epistemic actions correspond to public announcements of propositional facts, chains and trees to certain forms of private announcements, and graphs capture any propositional epistemic action. Possible applications of such planning fragments could e.g. be planning in games like Clue/Cluedo where actions can be seen as purely epistemic; or synthesis of protocols for secure communication (where


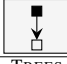

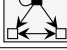
Underlying graphs of actions	Types of epistemic actions		
	Non-factual, propositional preconditions	Factual, propositional preconditions	Factual, epistemic preconditions
SINGLETONS 	NP-complete (Theorem 5.1)	PSPACE-hard [Jensen, 2014]	PSPACE-hard [Jensen, 2014]
CHAINS 	NP-complete (Theorem 5.2)	? (open question)	? (open question)
TREES 	PSPACE-complete (Theorem 5.3, Theorem 5.4)	? (open question)	? (open question)
GRAPHS 	in EXSPACE (Theorem 5.8)	in NON-ELEMENTARY [Yu <i>et al.</i> , 2013]	Undecidable [Bolander and Andersen, 2011]

Table 1: Complexity results for the plan existence problem.

the goal specifies who are allowed to know what).

Technically our fragments increase in complexity as we loosen the restrictions put on the underlying graph. Some planning tasks might therefore be simplified during a preprocessing phase so that a better upper bound can be guaranteed. As a case in point, we can preprocess any planning task and replace each graph action with a tree action that is equivalent up to a predetermined modal depth k (by unravelling). Letting k be the modal depth of the goal formula we obtain an equivalent planning task that can be solved using at most space polynomial in k and the size of the planning task. In automated planning such preprocessing is often used to achieve scalable planning systems.

Our results also allow us to prove that the plan verification problem (a subproblem of DEL model checking) is PSPACE-hard, even without non-deterministic union, thereby improving the result of [Aucher and Schwarzentruber, 2013].

Sections 2 and 3 present the core notions from epistemic planning. In Section 4 we improve the known upper bounds on the stabilisation of epistemic actions. This is put to use in Section 5 where we show novel results on the complexity of the plan existence problem. Section 6 presents our improvement to the plan verification problem, before we conclude and discuss future work in Section 7.

2 Background on Epistemic Planning

For the remainder of the paper we fix both an infinitely countable set of atomic propositions P and a finite set of agents Ag .

2.1 Dynamic epistemic logic

Definition 2.1 (Epistemic models and states). An *epistemic model* is a triple $M = (W, R, V)$ where the *domain* W is a non-empty set of worlds; $R : Ag \rightarrow 2^{W \times W}$ assigns an *epistemic (accessibility) relation* to each agent; and $V : P \rightarrow 2^W$ assigns a *valuation* to each atomic proposition. We write R_a for $R(a)$ and $wR_a v$ for $(w, v) \in R_a$. We often write W^M for W , R_a^M for R_a and V^M for V . For $w \in W$, the pair (M, w) is called an *epistemic state* whose *actual world* is w . (M, w) is *finite* when W is finite. Epistemic states are typically denoted by symbols such as s and s_0 .

The language of propositional logic over P is referred to as L_{Prop} , or sometimes simply the *propositional language*.

Definition 2.2 (Propositional action models and epistemic actions). A *propositional action model* is a triple $A = (E, Q, pre)$ where E is a non-empty and finite set of *events* called the *domain* of A ; $Q : Ag \rightarrow 2^{E \times E}$ assigns an *epistemic (accessibility) relation* to each agent; and $pre : E \rightarrow L_{\text{Prop}}$ assigns a *precondition* of the propositional language to each event. We write Q_a for $Q(a)$ and $eQ_a f$ for $(e, f) \in Q_a$. We often write E^A for E , Q_a^A for Q_a and pre^A for pre . For $e \in E$, the pair (A, e) is called an *epistemic action* whose *actual event* is e . Epistemic actions are typically denoted $\alpha, \alpha', \alpha_1$, etc.

Propositional action models are defined to fit exactly our line of investigation here, though other presentations consider preconditions of more complex languages and postconditions that allow for factual (ontic) change [Bolander and Andersen, 2011; Yu *et al.*, 2013].

The *dynamic language* L_D is generated by the BNF:

$$\varphi := p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box_a\varphi \mid \langle\alpha\rangle\varphi$$

where $a \in Ag$, $p \in P$ and α is an epistemic action. Here \Box_a denotes the *knowledge (or, belief) modality* where $\Box_a\varphi$ reads as “ a knows (or, believes) φ ”, and $\langle\alpha\rangle$ is the *dynamic modality* where $\langle\alpha\rangle\varphi$ reads as “ α is applicable and φ holds after executing α ”. The *epistemic language* L_E is the sublanguage of L_D that does not contain the dynamic modality. As usual we use $\Diamond_a\varphi := \neg\Box_a\neg\varphi$, and define by abbreviation \top , \perp and the boolean connectives \vee , \rightarrow , \leftrightarrow . Lastly, we define $\langle\alpha\rangle^0\varphi := \varphi$ and $\langle\alpha\rangle^k\varphi := \langle\alpha\rangle(\langle\alpha\rangle^{k-1}\varphi)$ for $k > 0$.

Definition 2.3 (Semantics). Let (M, w) be an epistemic state where $M = (W, R, V)$. For $a \in Ag$, $p \in P$ and $\varphi, \varphi' \in L_D$ we inductively define truth of formulas as follows, omitting the propositional cases:

$$\begin{aligned} (M, w) \models \Box_a\varphi & \quad \text{iff } (M, v) \models \varphi \text{ for all } wR_a v \\ (M, w) \models \langle\alpha\rangle\varphi & \quad \text{iff } (M, w) \models pre(e) \\ & \quad \text{and } (M \otimes A, (w, e)) \models \varphi \end{aligned}$$

where $\alpha = (A, e)$ is an epistemic action s.t. $A = (E, Q, pre)$, and the epistemic model $M \otimes A = (W', R', V')$ is defined via the *product update operator* \otimes by:

$$\begin{aligned} W' &= \{(v, f) \in W \times E \mid (M, v) \models pre(f)\}, \\ R'_a &= \{((v, f), (u, g)) \in W' \times W' \mid vR_a u, fQ_a g\}, \\ V'(p) &= \{(v, f) \in W' \mid v \in V(p)\} \text{ for } p \in P. \end{aligned}$$

For any epistemic state $s = (M, w)$ and epistemic action $\alpha = (A, e)$ satisfying $(M, w) \models pre^A(e)$, we define $s \otimes \alpha = (M \otimes A, (w, e))$. The epistemic state $s \otimes \alpha$ represents the result of executing α in s . Note that we have $s \models \langle\alpha\rangle\varphi$ iff $(M, w) \models pre^A(e)$ and $s \otimes \alpha \models \varphi$. Two formulas φ, φ' of L_D are called *equivalent* (written as $\varphi \equiv \varphi'$) when $s \models \varphi$ iff $s \models \varphi'$ for every epistemic state s .

Example 2.4. Consider the epistemic state s_1 of Figure 1. It represents a situation where p holds in the actual world (w), but where the two agents, a and b , don't know this: $s_1 \models p \wedge \neg\Box_a p \wedge \neg\Box_b p$. Consider now the epistemic action $\alpha_1 = (A, e)$ of the same figure. It represents a private announcement of p to agent a , that is, agent a is told that p holds (the actual event, e), but agent b thinks that nothing is




Figure 1: (Top left) An epistemic state s_1 . We mark each world (circle) with its name and the atomic propositions that are true. The actual world is coloured black. Edges show epistemic relations of the agents. (Bottom left) An epistemic action α_1 . We use the same conventions as for epistemic states, except an event (square) is marked by its name and its precondition. (Right) The epistemic model to the right is the result of execution of α_1 in s_1 , that is, $s_1 \otimes \alpha_1$.

happening (event f). The dynamic modality allows us to reason about the result of executing α_1 in s_1 , so for instance we have $s_1 \models \langle \alpha_1 \rangle (\Box_a p \wedge \neg \Box_b p \wedge \neg \Box_b \Box_a p)$: After agent a has been privately informed about p , she will know it, but b will still not know p , and will believe that a doesn't either. This fact can be verified by observing that $\Box_a p \wedge \neg \Box_b p \wedge \neg \Box_b \Box_a p$ is true in the epistemic state $s_1 \otimes \alpha_1$ of Figure 1.

2.2 Plan existence problem

Definition 2.5 (Planning tasks). An (epistemic) planning task is a triple $T = (s_0, \mathcal{L}, \varphi_g)$ where s_0 is a finite epistemic state called the *initial state*, \mathcal{L} is a finite set of epistemic actions called the *action library* and $\varphi_g \in L_E$ is called the *goal*. A *plan* for T is a finite sequence $\alpha_1, \dots, \alpha_j$ of epistemic actions from \mathcal{L} s.t. $s_0 \models \langle \alpha_1 \rangle \dots \langle \alpha_j \rangle \varphi_g$. The sequence $\alpha_1, \dots, \alpha_j$ can contain any number of repetitions, and can also be empty. We say that T is *solvable* if there exists a plan for T . The size of a planning task $T = (s_0, \mathcal{L}, \varphi_g)$ is given as follows. Following [Aucher and Schwarzen-truber, 2013], for any $\alpha = (A, e)$ in \mathcal{L} we define $|\alpha| = |Ag| \cdot |E^A|^2 + \sum_{e \in E} |pre(e)|$ as the *size of α* , where $|pre(e)|$ denotes the length of the (propositional) formula $pre(e)$. The size of an epistemic action is always a finite number, since the domain of any propositional action model and Ag are both finite. Let $P' \subseteq P$ be the finite set of atomic propositions that occur either in some precondition of an $\alpha \in \mathcal{L}$ or in φ_g . The *size T* is then $|T| = |P'| \cdot |Ag| \cdot |W^M|^2 + \sum_{\alpha \in \mathcal{L}} |\alpha| + |\varphi_g|$ where $s_0 = (M, w)$.

Note that a plan is nothing more than a sequence of epistemic actions leading to a goal. It is not hard to show that this definition is equivalent to the definition of a solution [Aucher and Bolander, 2013] and an explanatory diagnosis [Yu *et al.*, 2013], which are both special cases of a solution to a classical planning task as defined in [Ghallab *et al.*, 2004] (for the relation to classical planning tasks, see [Aucher and Bolander, 2013]).

Example 2.6. Consider again Figure 1. We'll use α_2 to refer to the private announcement of p to b , obtained simply by swapping the epistemic relations of a and b in α_1 . Consider the planning task $T = (s, \{\alpha_1, \alpha_2\}, \varphi_g)$ with $\varphi_g = \Box_a p \wedge$

$\Box_b p \wedge \neg \Box_a \Box_b p \wedge \neg \Box_b \Box_a p$. It is a planning task in which the only available actions are private announcements of p to either a or b , and the goal is for both a and b to know p , but without knowing that the other knows. A plan for T is α_1, α_2 , since $s \models \langle \alpha_1 \rangle \langle \alpha_2 \rangle \varphi_g$. In other words, first announcing p privately to a and then privately to b will achieve the goal of them both knowing p without knowing that each other knows.

Definition 2.7 (Plan existence problem). Let X denote a class of planning tasks. The *plan existence problem for X* , called $PLANEX(X)$ is the following decision problem: Given a planning task $T \in X$, does there exists a plan for T ?

3 Background on Iterating Epistemic Actions

To get to grips on the plan existence problem, we now consider the result of iterating a single epistemic action. We then proceed to derive a useful characterisation of exactly when a planning task is solvable.

Definition 3.1 (n -ary product). Let $\alpha = (A, e)$ be an epistemic action where $A = (E, Q, pre)$. We denote by $A^n = (E^n, Q^n, pre^n)$ the *n -ary product of A* . We define $E^0 = \{e\}$, $eQ_a^0 e$ for each $a \in Ag$, and $pre^0(e) = \top$. For $n > 0$ we define

- $E^n = \{(e_1, \dots, e_n) \mid e_i \in E \text{ for all } i = 1, \dots, n\}$,
- $Q_a^n = \{((e_1, \dots, e_n), (f_1, \dots, f_n)) \mid e_i Q_a f_i \text{ for all } i = 1, \dots, n\}$ for each $a \in Ag$, and
- $pre^n((e_1, \dots, e_n)) = \bigwedge_{i=1, \dots, n} pre(e_i)$.

The *n -ary product of α* is defined as $\alpha^n = (A^n, e^n)$, where e^n denotes $\underbrace{(e, e, \dots, e)}_n$.

This is not the standard definition of the n -ary product of an action model, which instead goes via a definition of the product update operator on action models. Definition 3.1 is equivalent to the standard definition when preconditions are of L_{Prop} . The following lemma is derived from the axiomatization of [Baltag *et al.*, 1998] (relying in particular on action composition), and is here stated for the case of the n -ary product and utilising that preconditions are of L_{Prop} .

Lemma 3.2. For any epistemic action α and any $\varphi \in L_E$ we have that $\langle \alpha \rangle^n \varphi \equiv \langle \alpha^n \rangle \varphi$.

This lemma expresses that executing an epistemic action n times is equivalent to executing its n -ary product once.

3.1 Bisimilarity and Stabilisation

Concerning n -ary products of epistemic actions, an interesting case is when executing the n -ary product is equivalent to executing the $(n + 1)$ -ary product. This puts an upper bound on the number of times the action needs to occur in a plan since epistemic actions with propositional preconditions commute [Löwe *et al.*, 2011]. To analyse this, we introduce notions of bisimulation and n -bisimulation on action models (slightly reformulated from [Sadzik, 2006]).

Definition 3.3 (Bisimilarity). Two epistemic actions $\alpha = (A, e)$ and $\alpha' = (A', e')$ are called *bisimilar*, written $\alpha \dot{\sim} \alpha'$, if there exists a (bisimulation) relation $Z \subseteq E^A \times E^{A'}$ containing (e, e') and satisfying for every $a \in Ag$:

- **[atom]** If $(f, f') \in Z$ then $pre^A(f) \equiv pre^{A'}(f')$,
- **[forth]** If $(f, f') \in Z$ and $fQ_a^A g$ then there is a $g' \in E^{A'}$ such that $f'Q_a^{A'} g'$ and $(g, g') \in Z$, and
- **[back]** If $(f, f') \in Z$ and $f'Q_a^{A'} g'$ then there is a $g \in E^A$ such that $fQ_a^A g$ and $(g, g') \in Z$.

Definition 3.4 (n -bisimilarity). Let $\alpha = (A, e)$ and $\alpha' = (A', e')$ be epistemic actions. They are 0 -bisimilar, written $\alpha \stackrel{\leftrightarrow}{\sim}_0 \alpha'$, if $pre^A(e) \equiv pre^{A'}(e')$. For $n > 0$, they are n -bisimilar, written $\alpha \stackrel{\leftrightarrow}{\sim}_n \alpha'$, if for every $a \in Ag$:

- **[atom]** $pre^A(e) \equiv pre^{A'}(e')$,
- **[forth]** If $eQ_a^A f$ then there is an $f' \in E^{A'}$ such that $e'Q_a^{A'} f'$ and $(A, f) \stackrel{\leftrightarrow}{\sim}_{n-1}(A', f')$, and
- **[back]** If $e'Q_a^{A'} f'$ then there is an $f \in E^A$ such that $eQ_a^A f$ and $(A, f) \stackrel{\leftrightarrow}{\sim}_{n-1}(A', f')$.

The modal depth $md(\varphi)$ of a formula φ is defined as: $md(p) = 0$; $md(\neg\varphi) = md(\varphi)$; $md(\varphi \wedge \psi) = \max\{md(\varphi), md(\psi)\}$; $md(\Box_a\varphi) = 1 + md(\varphi)$; $md(\langle\alpha\rangle\varphi) = md(\varphi)$. As epistemic actions have only propositional preconditions, $\langle\alpha\rangle$ -operators do not count towards the modal depth. This definition of modal depth, Lemma 3.5 and Definition 3.6 are all due to [Sadzik, 2006] (slightly reformulated).


Lemma 3.5. Let α, α' be two epistemic actions and $\varphi \in L_D$.

- 1) If $\alpha \stackrel{\leftrightarrow}{\sim} \alpha'$, then $\langle\alpha\rangle\varphi \equiv \langle\alpha'\rangle\varphi$.
- 2) If $md(\varphi) \leq n$ and $\alpha \stackrel{\leftrightarrow}{\sim}_n \alpha'$, then $\langle\alpha\rangle\varphi \equiv \langle\alpha'\rangle\varphi$.

Definition 3.6 (Stabilisation). Let α be an epistemic action.

- 1) α is $\stackrel{\leftrightarrow}{\sim}$ -stabilising at stage i if $\alpha^i \stackrel{\leftrightarrow}{\sim} \alpha^{i+k}$ for all $k \geq 0$.
- 2) α is $\stackrel{\leftrightarrow}{\sim}_n$ -stabilising at stage i if $\alpha^i \stackrel{\leftrightarrow}{\sim}_n \alpha^{i+k}$ for all $k \geq 0$.

Example 3.7. The 2-ary product α_1^2 of α_1 of Figure 1 is:



It is easy to check that $\alpha_1 \stackrel{\leftrightarrow}{\sim} \alpha_1^2$, using $Z = \{(e, (e, e)), (f, (f, f))\}$. This argument can be extended to show that α_1 is indeed $\stackrel{\leftrightarrow}{\sim}$ -stabilising at stage 1.

Since any epistemic action is finite, we have:

Lemma 3.8. If two epistemic actions are n -bisimilar for all n , then they are bisimilar.

3.2 Bounding the Number of Iterations

We're now ready to present our characterisation of when a planning task is solvable. We note that Proposition 3.9 below echoes the sentiment of [Yu *et al.*, 2013, Theorem 5.15], in that it states the conditions under which we can restrict the search space when looking for a plan.

Proposition 3.9. Let $T = (s_0, \{\alpha_1, \dots, \alpha_m\}, \varphi_g)$ be a planning task and $B \in \mathbb{N}$. Suppose one of the following holds:

- 1) Every α_i is $\stackrel{\leftrightarrow}{\sim}$ -stabilising at stage B , or

2) $md(\varphi_g) = n$ and every α_i is $\stackrel{\leftrightarrow}{\sim}_n$ -stabilising at stage B . Then T is solvable iff there exists $k_1, \dots, k_m \leq B$ s.t. $s_0 \models \langle\alpha_1^{k_1}\rangle \dots \langle\alpha_m^{k_m}\rangle\varphi_g$.

Proof. Assume 2) holds (the case of 1) is similar). Assume T is solvable, and let $\alpha_{i_1}, \dots, \alpha_{i_j}$ be a plan for T . Due to commutativity of propositional action models so is any permutation of $\alpha_{i_1}, \dots, \alpha_{i_j}$ [Yu *et al.*, 2013]. We therefore have $s_0 \models \langle\alpha_1\rangle^{k'_1} \dots \langle\alpha_m\rangle^{k'_m}\varphi_g$ for some choice of $k'_i \geq 0$. Using Lemma 3.2, it follows that $s_0 \models \langle\alpha_1^{k'_1}\rangle \dots \langle\alpha_m^{k'_m}\rangle\varphi_g$. We now let $k_i = \min(k'_i, B)$ for all i . By assumption, $md(\varphi_g) = n$ and so by definition $md(\langle\alpha\rangle\varphi_g) = n$ for any epistemic action α . Combining this with the assumption that every α_i is $\stackrel{\leftrightarrow}{\sim}_n$ -stabilising at stage $B \geq k_i$, we apply 2) of Lemma 3.5 m times to conclude that $s_0 \models \langle\alpha_1^{k_1}\rangle \dots \langle\alpha_m^{k_m}\rangle\varphi_g$, as required. The proof of the other direction follows readily from Lemma 3.2 and the definition of $\langle\alpha\rangle^k$. \square

Let $T = \{s_0, \mathcal{L}, \varphi_g\}$ be a planning task with $md(\varphi_g) = n$. Given the proposition above, to show that T is solvable we only need to find the correct number of times to iterate each of the actions in \mathcal{L} , and these numbers never have to exceed B for actions that are $\stackrel{\leftrightarrow}{\sim}_n$ -stabilising at stage B . The following result, due to [Sadzik, 2006], shows that such a bound B exists for any epistemic action.

Lemma 3.10. Let $\alpha = (A, e)$ be an epistemic action and n a natural number. Then α is $\stackrel{\leftrightarrow}{\sim}_n$ -stabilising at stage $|E^A|^n$.

4 Better Bounds for Action Stabilisation

In this section, we prove an original contribution, Lemma 4.2, that generalises Sadzik's Lemma 3.10 by giving a better bound for action stabilisation. The overall point is this: Sadzik gets an unnecessarily high upper bound on when an epistemic action (A, e) stabilises by considering it possible that any event can have up to $|E^A|$ successors. We get a better bound by counting paths.

Definition 4.1 (Underlying graphs). Let (A, e) be an epistemic action. We define $\mathbf{Q}^A = \cup_{a \in Ag} Q_a^A$. The *underlying graph* of (A, e) is the directed graph (A, \mathbf{Q}^A) with root e .

Let (A, e) denote an epistemic action. Note that $(e, f) \in \mathbf{Q}^A$ iff there is an edge from e to f in A labelled by *some* agent. Standard graph-theoretical notions carry over to epistemic actions via their underlying graphs. For instance, we define a *path of length n* in (A, e) as a path of length n in the underlying graph, that is, a sequence $(e_1, e_2, \dots, e_{n+1})$ of events such that $(e_i, e_{i+1}) \in \mathbf{Q}^A$ for all $i = 1, \dots, n$ (we allow $n = 0$ and hence paths of length 0). A *path of length $\leq n$* is a path of length at most n . A *maximal path of length $\leq n$* is a path of length $\leq n$ that is not a strict prefix of any other path of length $\leq n$. We use $\text{mpaths}_n(e)$ to denote the number of distinct maximal paths of length $\leq n$ rooted at e . If all nodes have successors, this number is simply the number of distinct paths of length n . Note that $\text{mpaths}_n(e)$ is always a positive number, as there is always at least one path rooted at e (even if e has no outgoing edges, there is still the path of length 0). Note also that for any $n > 0$ and any event e having at least one successor in the underlying graph:

$$\text{mpaths}_n(e) = \sum_{e \in Q^A f} \text{mpaths}_{n-1}(f). \quad (1)$$

In the epistemic action α_1 of Figure 1 we have $\text{mpaths}_2(e) = 3$, since there are three paths of length 2, (e, e, e) , (e, e, f) and (e, f, f) , and no shorter maximal paths.

Lemma 4.2. *Let $\alpha = (A, e_0)$ be an epistemic action and n any natural number. Then α is \Leftrightarrow_n -stabilising at stage $\text{mpaths}_n(e_0)$.*

Proof. When $\mathbf{f} = (f_1, \dots, f_m) \in E^{A^m}$ and $e \in A$, we use $\text{occ}(e, \mathbf{f})$ to denote the number of occurrences of e in f_1, \dots, f_m . For instance we have $\text{occ}(e, (e, e, f, f)) = 2$. We now prove the following property $\mathcal{P}(n)$ by induction on n .

$\mathcal{P}(n)$: If $\mathbf{e} \in E^{A^{k+1}}$ and $\mathbf{e}' \in E^{A^k}$ only differ by some event e^* occurring at least $\text{mpaths}_n(e^*) + 1$ times in \mathbf{e} and at least $\text{mpaths}_n(e^*)$ times in \mathbf{e}' , then $(A^{k+1}, \mathbf{e}) \Leftrightarrow_n (A^k, \mathbf{e}')$.

Base case $\mathcal{P}(0)$: Since $\text{mpaths}_0(e^*) = 1$, \mathbf{e} and \mathbf{e}' as described above must contain exactly the same events (but not necessarily with the same number of occurrences). By definition of the n -ary product of an epistemic action we get $\text{pre}^{A^{k+1}}(\mathbf{e}) \equiv \text{pre}^{A^k}(\mathbf{e}')$. This shows $(A^{k+1}, \mathbf{e}) \Leftrightarrow_0 (A^k, \mathbf{e}')$. For the induction step, assume that $\mathcal{P}(n-1)$ holds. Given \mathbf{e} and \mathbf{e}' as described in $\mathcal{P}(n)$, we need to show $(A^{k+1}, \mathbf{e}) \Leftrightarrow_n (A^k, \mathbf{e}')$. **[atom]** is proved as $\mathcal{P}(0)$.

[forth]: Let a and \mathbf{f} be chosen such that $e Q_a^{A^{k+1}} \mathbf{f}$. We need to find \mathbf{f}' such that $e' Q_a^{A^k} \mathbf{f}'$ and $(A^{k+1}, \mathbf{f}) \Leftrightarrow_{n-1} (A^k, \mathbf{f}')$.

Claim. There exists an f^* such that $e^* Q_a^A f^*$ and $\text{occ}(f^*, \mathbf{f}) \geq \text{mpaths}_{n-1}(f^*) + 1$.

Proof of Claim. By contradiction: Suppose $\text{occ}(f, \mathbf{f}) \leq \text{mpaths}_{n-1}(f)$ for all f with $e^* Q_a^A f$. Since $e Q_a^{A^{k+1}} \mathbf{f}$, the number of occurrences of e^* in \mathbf{e} is equal or less than the number of occurrences of Q_a -successors of e^* in \mathbf{f} . Hence we get

$$\begin{aligned} \text{occ}(e^*, \mathbf{e}) &\leq \sum_{e^* Q_a^A f} \text{occ}(f, \mathbf{f}) \\ &\leq \sum_{e^* Q_a^A f} \text{mpaths}_{n-1}(f) \quad (\text{by assumption}) \\ &\leq \sum_{e^* Q_a^A f} \text{mpaths}_{n-1}(f) \quad (\text{by } \mathbf{Q}^A = \bigcup_{a \in Ag} Q_a^A) \\ &= \text{mpaths}_n(e^*) \quad (\text{by equation (1)}). \end{aligned}$$

However, this directly contradicts the assumption that e^* occurs at least $\text{mpaths}_n(e^*) + 1$ times in \mathbf{e} , and hence the proof of the claim is complete.

Let f^* be as guaranteed by the claim. Now we build \mathbf{f}' to be exactly like \mathbf{f} , except we omit one of the occurrences of f^* (we do not have to worry about the order of the elements of the vectors, since any two vectors only differing in order are bisimilar [Sadzik, 2006]). Since \mathbf{f} and \mathbf{f}' now only differ in f^* occurring at least $\text{mpaths}_{n-1}(f^*) + 1$ times in \mathbf{f} and at least $\text{mpaths}_{n-1}(f^*)$ times in \mathbf{f}' , we can use the induction hypothesis $\mathcal{P}(n-1)$ to conclude that $(A^{k+1}, \mathbf{f}) \Leftrightarrow_{n-1} (A^k, \mathbf{f}')$, as required. **[back]:** This is the easy direction and is omitted.

Now we have proved $\mathcal{P}(n)$ for all n . Given n , from $\mathcal{P}(n)$ it follows that $(A^{k+1}, e_0^{k+1}) \Leftrightarrow_n (A^k, e_0^k)$ for all $k \geq \text{mpaths}_n(e_0)$. And from this it immediately follows that (A, e_0) is \Leftrightarrow_n -stabilising at stage $\text{mpaths}_n(e_0)$. \square




Figure 2: Initial state and actions used in Theorem 5.1.

procedure PlanExists $((s_0, \{\alpha_1, \dots, \alpha_m\}, \varphi_g), B)$
 | a) Guess a vector $(k_1, \dots, k_m) \in \{0, \dots, B\}^m$.
 | b) Accept when $s_0 \models \langle \alpha_1 \rangle^{k_1} \dots \langle \alpha_m \rangle^{k_m} \varphi_g$.

Figure 3: Non-deterministic algorithm for the plan existence problem.

5 Complexity of the Plan Existence Problem

5.1 Singleton and Chain Epistemic Actions

We define SINGLETONS as the class of planning tasks $(s_0, \mathcal{L}, \varphi_g)$, where every $\alpha = (A, e)$ in \mathcal{L} is a *singleton*; i.e. E^A contains a single event.

Theorem 5.1. $\text{PLANEX}(\text{SINGLETONS})$ is NP-complete.

Proof. For any singleton epistemic action there is at most *one* maximal path of length $\leq n$ for all n . Hence, by Lemma 4.2 and 3.8, such actions are \Leftrightarrow -stabilising at stage 1. **IN NP:** Follows from Theorem 5.2 below as SINGLETONS is contained in CHAINS. **NP-hard:** We give a polynomial-time reduction from SAT. Let $\varphi(p_1, \dots, p_m)$ be a propositional formula where p_1, \dots, p_m are the atomic propositions in φ . We construct $T = (s_0, \{\alpha_1, \dots, \alpha_m\}, \varphi_g)$ s.t. s_0 and each α_i are as in Figure 2 and $\varphi_g = \varphi(\diamond_a p_1, \dots, \diamond_a p_m)$ is the formula φ in which each occurrence of p_i is replaced by $\diamond_a p_i$. For any propositional valuation ν , let s_ν be the restriction of s_0 s.t. there is an a -edge from r to w_i in s_ν iff $\nu \models p_i$. This means $\nu \models p_i$ iff $s_\nu \models \diamond_a p_i$, and so from our construction of φ_g we have $\nu \models \varphi$ iff $s_\nu \models \varphi_g$. Observe now that $s_0 \otimes \alpha_i$ is exactly the restriction of s_0 so that there is no a -edge from r to w_i , and that α_i is the only action affecting this edge. Let $\bar{\nu}(p_i) = 0$ if $\nu \models p_i$ and $\bar{\nu}(p_i) = 1$ otherwise. We now have that φ is satisfiable iff there is a ν s.t. $\nu \models \varphi$ iff $s_\nu \models \varphi_g$ iff $s_0 \models \langle \alpha_1 \rangle^{\bar{\nu}(p_1)} \dots \langle \alpha_m \rangle^{\bar{\nu}(p_m)} \varphi_g$ iff T is solvable, where the last equivalence follows from Proposition 3.9 since $\bar{\nu}(p_i) \in \{0, 1\}$ and each α_i is \Leftrightarrow -stabilising at stage 1. This shows that φ is satisfiable iff T is solvable. \square

For an epistemic action $\alpha = (A, e)$ we say that α is a *chain* if its underlying graph (A, \mathbf{Q}^A) is a 1-ary tree whose unique leaf may be \mathbf{Q}^A -reflexive. We define CHAINS as the class of planning tasks $(s_0, \mathcal{L}, \varphi_g)$ where every epistemic action in \mathcal{L} is a chain.

Theorem 5.2. $\text{PLANEX}(\text{CHAINS})$ is NP-complete.

Proof. **IN NP:** For any chain epistemic action there is at most *one* maximal path of length $\leq n$ for all n , hence any such action is \Leftrightarrow -stabilising at stage 1 using Lemmas 4.2 and 3.8. It therefore follows from Proposition 3.9 that, for any $T \in \text{CHAINS}$, $\text{PlanExists}(T, 1)$ of Figure 3 is accepting iff T is solvable. We must show step b) to run in polynomial

time. Now if α is a chain and s an epistemic state, then the number of worlds reachable from the actual world in $s \otimes \alpha$ is at most the number of worlds in s . By only keeping the reachable worlds after each successive product update, we get the required, as the goal is in L_E .¹ **NP-hard:** Follows from Theorem 5.1 as SINGLETONS is contained in CHAINS. \square

5.2 Tree Epistemic Actions

We now turn to epistemic actions whose underlying graph is a any tree. Formally, an epistemic action (A, e) is called a *tree* when the underlying graph (A, \mathbf{Q}^A) is a tree whose leaves may be \mathbf{Q}^A -reflexive. We call TREES the class of planning tasks $(s_0, \mathcal{L}, \varphi_g)$ where all epistemic actions in \mathcal{L} are trees.

Theorem 5.3. PLANEX(TREES) is in PSPACE.

Proof. Consider any tree action $\alpha = (A, e)$ and let $l(\alpha)$ denote its number of leaves. As α is a tree, we get $\text{mpaths}_n(e) \leq l(\alpha)$ for any n . Using Lemma 4.2 and 3.8, any tree epistemic action α is \Leftrightarrow -stabilising at stage $l(\alpha)$. From Proposition 3.9 we therefore have, for any $T \in \text{TREES}$, that $\text{PlanExists}(T, \max(l(\alpha_1), \dots, l(\alpha_m)))$ of Figure 3 is accepting iff T is solvable. Step b) can be done in space polynomial in the size of the input [Aucher and Schwarzentruber, 2013]. Hence, the plan existence problem for TREES is in NPSpace and therefore in PSPACE by Savitch's Thm. \square

We now sketch a proof of PSPACE-hardness of PLANEX(TREES), by giving a polynomial-time reduction from the PSPACE-hard problem QSAT (satisfiability of quantified boolean formulas) to PLANEX(TREES). For any quantified boolean formula $\Phi = Q_1 p_1 \dots Q_n p_n \varphi[p_1, \dots, p_n]$ with $Q_i \in \{\forall, \exists\}$, we define the planning task $T_\Phi = (s_0, \{\alpha_1, \dots, \alpha_n\}, \varphi_{\text{sat}} \wedge \varphi_{\text{all}})$ where s_0 and each α_i are as in Figure 4 (every edge implicitly labelled by a),

$$\varphi_{\text{sat}} = O_1 \dots O_n \varphi[\diamond_a^1 \square_a \perp, \dots, \diamond_a^n \square_a \perp], \text{ and}$$

$$\varphi_{\text{all}} = \diamond_a^{n+1} \square_a \perp \wedge \dots \wedge \diamond_a^{2n} \square_a \perp,$$

where $O_i = \diamond_a$ if $Q_i = \exists$ and $O_i = \square_a$ if $Q_i = \forall$. Then $|T_\Phi|$ is polynomial in $|\Phi|$ and $T_\Phi \in \text{TREES}$. By Lemmas 5.6 and 5.7 below we get T_Φ is solvable iff Φ is true. Hence:

Theorem 5.4. PLANEX(TREES) is PSPACE-hard.

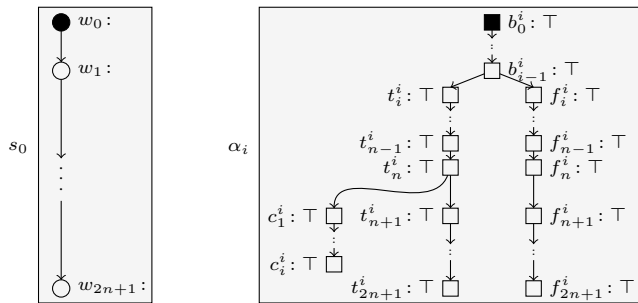


Figure 4: Initial state and actions used in Theorem 5.4.

¹Observe that even if each action in $\alpha_1, \dots, \alpha_m$ is \Leftrightarrow -stabilising at stage 1, this is not a sufficient condition for membership in NP as we must also be able to verify the plan in polynomial time.




Figure 5: Binary decision tree simulated by $s_0 \otimes \alpha_1 \otimes \alpha_2$ ($n = 2$).

The reduction is based on the idea that we can simulate a (complete) binary decision tree using $s' = s_0 \otimes \alpha_1 \otimes \dots \otimes \alpha_n$. Each world at depth n of s' simulates a valuation, using the convention that p_i is true iff there is a maximal chain of length i in this world. By nesting belief modalities we can check if such a chain exists. Each action α_i makes two copies of every node between depth i and n , which is how we can simulate every valuation.

A world \mathbf{w} at depth $i \leq n$ of s' is called an *i-world*. It can now be verified that any *i-world* is of the form $(w_i, v_i^1, \dots, v_i^i, b_i^{i+1}, \dots, b_i^n)$ where $v_i^j \in \{t_i^j, f_i^j\}$. See also Figure 5. For any *i-world* \mathbf{w} , we define a propositional valuation $\nu_{\mathbf{w}}$ on $\{p_1, \dots, p_i\}$ by $\nu_{\mathbf{w}} \models p_j$ iff t_i^j occurs in \mathbf{w} . We use $\mathbf{w}_0 = (w_0, b_0^1, \dots, b_0^n)$ to denote the single 0-world in s' (the actual world of s'), and define M' so that $s' = (M', \mathbf{w}_0)$.

Lemma 5.5. Let \mathbf{w} be any n -world. Then $(M', \mathbf{w}) \models \varphi[\diamond_a^1 \square_a \perp, \dots, \diamond_a^n \square_a \perp]$ iff $\nu_{\mathbf{w}} \models \varphi[p_1, \dots, p_n]$ is true.

Proof sketch. Due to the c_1^i, \dots, c_i^i chain in each α_i , we have for any n -world \mathbf{w} and $i \leq n$ that $(M', \mathbf{w}) \models \diamond_a^i \square_a \perp$ iff t_i^n occurs in \mathbf{w} , from which the result readily follows. \square

We say that an n -world \mathbf{w} is *accepting* if $(M', \mathbf{w}) \models \varphi[\diamond_a^1 \square_a \perp, \dots, \diamond_a^n \square_a \perp]$, and for $i < n$ we say that the i -world \mathbf{w} is accepting if some (every) child \mathbf{w}' of \mathbf{w} is accepting and $O_i = \diamond_a$ ($O_i = \square_a$).

Lemma 5.6. T_Φ is solvable iff \mathbf{w}_0 is accepting.

Proof sketch. As acceptance for $i < n$ exactly corresponds to the $O_1 \dots O_n$ prefix, we use Lemma 5.5 to show that $(M', \mathbf{w}_0) \models \varphi_{\text{sat}}$ iff \mathbf{w}_0 is accepting. Now we must show: 1) $(M', \mathbf{w}_0) \models \varphi_{\text{all}}$, and then 2) T_Φ is solvable iff $\alpha_1, \dots, \alpha_n$ is plan for T_Φ . We omit proofs of both 1) and 2). \square

Lemma 5.7. Φ is true iff \mathbf{w}_0 is accepting.

Proof sketch. Let \mathbf{w} denote any i -world. Let $\bar{\nu}_{\mathbf{w}}(p_i) = \top$ if $\nu_{\mathbf{w}} \models p_i$ and $\bar{\nu}_{\mathbf{w}}(p_i) = \perp$ otherwise. We define $\Phi_{\mathbf{w}} = Q_{i+1} p_{i+1} \dots Q_n p_n \varphi[\bar{\nu}_{\mathbf{w}}(p_1), \dots, \bar{\nu}_{\mathbf{w}}(p_i), p_{i+1}, \dots, p_n]$.

By induction on k we now show: If $k \leq n$ and \mathbf{w} is an $(n - k)$ -world, then $\Phi_{\mathbf{w}}$ is true iff \mathbf{w} is accepting. For the base case, $k = 0$ and \mathbf{w} is an n -world, hence $\varphi[\bar{\nu}_{\mathbf{w}}(p_1), \dots, \bar{\nu}_{\mathbf{w}}(p_n)] (= \Phi_{\mathbf{w}})$ is true iff \mathbf{w} is accepting by Lemma 5.5. For the induction step we assume that for any $(n - (k - 1))$ -world \mathbf{w}' , $\Phi_{\mathbf{w}'}$ is true iff \mathbf{w}' is accepting. Let \mathbf{w} be an $(n - k)$ -world. By construction, \mathbf{w} has two children \mathbf{v} and \mathbf{u} . We can then show that $\Phi_{\mathbf{v}}$ and $\Phi_{\mathbf{u}}$ are as $\Phi_{\mathbf{w}}$, except the $Q_{n-k+1} p_{n-k+1}$ prefix and that one sets p_{n-k+2} true and the other sets p_{n-k+2} false. Thus $\Phi_{\mathbf{w}}$ is true iff $Q_{n-k+1} = \exists$ (or, $Q_{n-k+1} = \forall$) and $\Phi_{\mathbf{w}'}$ is true for some (every) child \mathbf{w}'

of w . Using the induction hypothesis, we get that Φ_w is true iff w' is accepting for some (every) child w' of w . Hence, Φ_w is true iff w is accepting, by definition. This concludes the induction proof. For $k = n$ it follows that Φ_{w_0} is true iff w_0 is accepting. Since $\Phi = \Phi_{w_0}$ we are done. \square

5.3 Arbitrary Epistemic Actions

We call **GRAPHS** the class of planning tasks $(s_0, \mathcal{L}, \varphi_g)$ where all event models in \mathcal{L} are arbitrary graphs. In this case, the original result by Sadzik (Lemma 3.10) is sufficient.

Theorem 5.8. **PLANEX(GRAPHS) is in EXPSPACE.**

Proof. We consider $(s_0, \{\alpha_1, \dots, \alpha_m\}, \varphi_g) \in \mathbf{GRAPHS}$ with $\alpha_i = (A_i, e_i)$ and $md(\varphi_g) = d$. By Lemma 3.10, each α_i is $\stackrel{\perp}{\leftrightarrow}_d$ -stabilising at stage $|E^{A_i}|^d$. It now follows from Proposition 3.9 that **PlanExists** $(T, \max\{|E^{A_1}|^d, \dots, |E^{A_m}|^d\})$ of Figure 3 is accepting iff T is solvable. The algorithm runs in $\mathbf{NEXPSPACE} = \mathbf{EXPSPACE}$. \square

6 Complexity of the Plan Verification Problem

The plan verification problem is defined as the following decision problem: Given a finite epistemic state s_0 and a formula of the form $\langle \alpha_1 \rangle \dots \langle \alpha_j \rangle \varphi_g$, does $s_0 \models \langle \alpha_1 \rangle \dots \langle \alpha_j \rangle \varphi_g$ hold? The plan verification problem can be seen as a restriction of the model checking problem in DEL. A similar reduction as for Theorem 5.4 gives that:

Theorem 6.1. *The plan verification problem (restricted to propositional action models that are trees) is PSPACE-hard.*

Model checking in DEL with the non-determinism operator \cup included in the language has already been proved PSPACE-hard [Aucher and Schwarzenrüber, 2013]. Theorem 6.1 implies that model checking in DEL is PSPACE-hard even without this operator. A similar result has been independently proved in [van de Pol *et al.*, 2015].

7 Future Work

We remind the reader that an overview of our contributions are found in Table 1 and proceed to discuss future work.

Since propositional STRIPS planning is PSPACE-complete [Bylander, 1994], efficient planning systems have used relaxed planning tasks in order to efficiently compute precise heuristics. For instance, the highly influential Fast-Forward planning system [Hoffmann and Nebel, 2001] relaxes planning tasks by ignoring delete lists. Our contributions here show that restrictions on the graphs underlying epistemic actions crucially affect computational complexity. This, in combination with restrictions on preconditions and postconditions (factual change), provides a platform for investigating (tractable) relaxations of epistemic planning tasks, and hence for the development of efficient epistemic planning systems.

References

- [Andersen *et al.*, 2012] Mikkel Birkegaard Andersen, Thomas Bolander, and Martin Holm Jensen. Conditional epistemic planning. In Luis Fariñas del Cerro, Andreas Herzig, and Jérôme Mengin, editors, *JELIA*, volume 7519 of *Lecture Notes in Computer Science*, pages 94–106. Springer, 2012.
- [Aucher and Bolander, 2013] Guillaume Aucher and Thomas Bolander. Undecidability in epistemic planning. In Francesca Rossi, editor, *IJCAI. IJCAI/AAAI*, 2013.
- [Aucher and Schwarzenrüber, 2013] Guillaume Aucher and François Schwarzenrüber. On the complexity of dynamic epistemic logic. In *TARK*, 2013.
- [Baltag *et al.*, 1998] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In *TARK*, 1998.
- [Bolander and Andersen, 2011] Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for single and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.
- [Bylander, 1994] Tom Bylander. The computational complexity of propositional strips planning. *Artificial Intelligence*, 69(1-2):165–204, 1994.
- [Ghallab *et al.*, 2004] Malik Ghallab, Dana S. Nau, and Paolo Traverso. *Automated Planning: Theory and Practice*. Morgan Kaufmann, 2004.
- [Hoffmann and Nebel, 2001] Jörg Hoffmann and Bernhard Nebel. The ff planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Resesearch*, 14(1):253–302, May 2001.
- [Jensen, 2014] Martin Holm Jensen. *Epistemic and Doxastic Planning*. PhD thesis, Technical University of Denmark, 2014. DTU Compute PHD-2014.
- [Löwe *et al.*, 2011] Benedikt Löwe, Eric Pacuit, and Andreas Witzel. Del planning and some tractable cases. In Hans van Ditmarsch, Jérôme Lang, and Shier Ju, editors, *Logic, Rationality, and Interaction*, volume 6953 of *Lecture Notes in Computer Science*, pages 179–192. Springer Berlin Heidelberg, 2011.
- [Maubert, 2014] Bastien Maubert. *Logical foundations of games with imperfect information: uniform strategies*. PhD thesis, 2014. PhD thesis directed by Pinchinat, Sophie and Aucher, Guillaume. Université de Rennes 1.
- [Miller and Moss, 2005] Joseph S Miller and Lawrence S Moss. The undecidability of iterated modal relativization. *Studia Logica*, 79(3):373–407, 2005.
- [Sadzik, 2006] Tomasz Sadzik. Exploring the Iterated Update Universe. ILLC Publications PP-2006-26, 2006.
- [van Benthem, 2003] Johan van Benthem. One is a lonely number: On the logic of communication. 2003.
- [van de Pol *et al.*, 2015] Iris van de Pol, Iris van Rooij, and Jakub Szymanik. Parameterized complexity results for a model of theory of mind based on dynamic epistemic logic. In *TARK*, 2015.
- [Yu *et al.*, 2013] Quan Yu, Ximing Wen, and Yongmei Liu. Multi-agent epistemic explanatory diagnosis via reasoning about actions. In Francesca Rossi, editor, *IJCAI. IJCAI/AAAI*, 2013.