

Dynamic Term-Modal Logics for First-Order Epistemic Planning

Andreas Achen¹,
Andrés Occhipinti Liberman²
and Rasmus K. Rendsvig³ (*corresponding*)

¹London School of Economics

²DTU Compute, Technical University of Denmark

³Center for Information and Bubble Studies, University of Copenhagen
Karen Blixens Plads 8, 2300 Copenhagen S, Denmark. rasmus@hum.ku.dk

Abstract

Many classical planning frameworks are built on first-order languages. The first-order expressive power is desirable for compactly representing actions via schemas, and for specifying quantified conditions such as $\neg\exists x \text{blocks_door}(x)$. In contrast, several recent epistemic planning frameworks are built on propositional epistemic logic. The epistemic language is useful to describe planning problems involving higher-order reasoning or epistemic goals such as $K_a \neg$ -problem.

This paper develops a first-order version of Dynamic Epistemic Logic (DEL). In this framework, for example, $\exists x K_x \exists y \text{blocks_door}(y)$ is a formula. The formalism combines the strengths of DEL (higher-order reasoning) with those of first-order logic (lifted representation) to model multi-agent epistemic planning. The paper introduces an epistemic language with a possible-worlds semantics, followed by novel dynamics given by first-order action models and their execution via product updates. Taking advantage of the first-order machinery, epistemic action schemas are defined to provide compact, problem-independent domain descriptions, in the spirit of PDDL.

Concerning metatheory, the paper defines axiomatic normal term-modal logics, shows a Canonical Model Theorem-like result which allows establishing completeness through frame characterization formulas, shows decidability for the finite agent case, and shows a general completeness result for the dynamic extension by reduction axioms.

Keywords: epistemic planning, planning formalisms, multi-agent systems, term-modal logic, dynamic epistemic logic

1 Introduction

Most classical planning languages are first-order. Standard formalisms like PDDL [59] and ADL [70], for example, are first-order. One major reason for using a first-order language over a propositional one is that variables can be used to describe actions compactly. For instance, in the PDDL description of BlocksWorld, the *action schema* $\text{stack}(X, Y)$ uses variables X and Y to represent generic blocks and state the preconditions and effects of all actions of the form: “put block X on top of block Y ”. This

is possible because the action of stacking block A on block B has the same type of effects as the action of stacking block C on D ; only the names of the blocks are different. Action schemas use variables to exploit this repeated structure in actions, resulting in action representations whose size is independent of the number of objects in a domain. While $\text{stack}(X, Y)$ describes the preconditions and effects of performing a stack action on any two blocks, regardless of total number of blocks, with a propositional language each stack action has to be represented by a separate model, yielding $n^2 - n$ propositional models for a domain with n blocks. Generally, given an action schema with k variables and n constant symbols standing for domain objects, the schema has up to n^k different instantiations, each requiring a separate model in a propositional representation.

Dynamic Epistemic Logic (DEL) has proved to be a very expressive framework for *epistemic planning*, i.e., planning explicitly involving e.g. knowledge or belief. DEL uses the language of propositional epistemic logic to describe the knowledge or belief held by a community of agents. This language is built from a set of propositional atoms, standard logical connectives, and modal operators K_i for each agent i in a fixed set of agent indices $I = \{1, \dots, n\}$. An example of a formula is $K_1 p \wedge K_2 K_1 p$, which expresses that agent 1 knows the propositional atom p and that agent 2 knows that agent 1 knows p . Actions in DEL are described by so-called *action models* [6, 7] or variants thereof. Action models describe preconditions and effects of events, and provide a rich framework for representing the agents' uncertainty about such events. However, as action models are based on the propositional epistemic language, propositional DEL cannot achieve the generality of action schemas. Variabilized, general descriptions are not possible, so one action model is required for each action. Hence, while propositional DEL adds a great deal of expressivity to planning, this comes at a cost in terms of representational succinctness.

This paper presents a DEL-based epistemic planning framework built on *epistemic term-modal logic*. The underlying language is first-order and includes modalities indexed by first-order terms. Examples of formulas include $K_c \text{on}(A, B)$ (agent c knows that block A is on block B), $K_c \exists x \text{on}(x, B)$ (c knows that there is a block on top of B) and $\forall y K_y \exists x \text{on}(x, B)$ (all agents know that there is a block on top of B). Term-modal languages thus extend the expressive power of first-order modal languages by treating modal operators *both* as operators *and* as predicates.

In addition to higher-order knowledge expressions, the first-order apparatus of epistemic term-modal logic allows for domain descriptions in terms of objects and relations, as well as abstract reasoning via variables and quantification. The term-modal aspect ensures that these first-order aspects also extend to agents and their knowledge. Importantly, the presence of variables enables the definition of *epistemic action schemas*. Epistemic action schemas can be exponentially more succinct than standard DEL event models (see Section 5.3). Moreover, epistemic action schemas provide an action representation whose size is independent of the number of agents and objects in the domain. We consider the development of this epistemic planning framework our first main contribution.

Our second main contribution is the development of term-modal logic, its dynamic extension and the metatheory for both. Many papers have been dedicated to term-modal logic and its metatheory (see Section 9.3 for a detailed review), but due to the many complications that may arise in such generalized first-order modal systems, no general completeness results have been shown. In this paper, we define a rich but well-behaved semantics that allow us to define axiomatic normal term-modal logics and show a Canonical Model Theorem-like result that allow completeness results through frame characterization formulas. Adding reduction axioms to the term-modal logics then allow us to show completeness for the dynamic extension.

The paper progresses as follows. Section 2 presents *SelectiveCommunication* used as running example of epistemic planning. Section 3 presents term-modal logical languages and Section 4 de-

finer state representations: first-order Kripke models where the agent set is a part of the domain of quantification. Section 5 introduces action representations (action models) and how these may be succinctly represented as action schemas. The action representations are used in Section 6 to define epistemic planning problems and related notions, and an example describing a term-modal planning domain and problem using a ‘PDDL-like syntax’ is there given. Section 7 details how to extend the term-modal language to allow reasoning about actions, Section 8 turns to axiomatic systems and metatheory, while Section 9 turns to related work on epistemic planning, dynamic epistemic logic and term-modal logic, respectively. Section 10 contains final remarks and open questions. All proofs may be found in Appendix A.

2 A Running Example

Throughout the paper, we illustrate the planning formalism with a simple running example in a variant of the *SelectiveCommunication* (SC) domain, adapted from [46]. Here we describe it informally, but it will serve as an example for the various formal notions throughout the paper. In the $SC(n, m, k, \ell)$ domain, there are n agents. Each agent is initially in one of m rooms arranged in a corridor. There are k boxes distributed in the rooms, each having one of ℓ available colors. See Figure 1 for an example. Agents can perform four types of actions:

- $Move(agent, room1, room2)$: agents can move from a room to a contiguous room, by going right or left. In this adaptation of the domain, we model the move actions as partially observable: if agent α is in room ρ_i and moves to room ρ_j , only the agents in either of the rooms can see that α ’s location has changed.
- $SenseLoc(agent, agent_or_box, room)$: while in a room, agents can sense the location of other agents or boxes in that room. Other agents in the room notice the sensing action.
- $SenseCol(agent, color, box, room)$: agents can sense the color of a box when they are in the same room as the box. Other agents in the room notice the sensing action.
- $Announce(agent, color, box, room)$: agents can make announcements concerning the colors of boxes. If α makes an announcement in a room, all agents in the same room or in a contiguous room will hear what was announced. α can use announcements to ensure that some agents get to know the truth value of some φ while the remaining agents do not.

A specific choice for the parameters n, m, k, ℓ yields an *instance* of the SC domain. For example, $SC(3, 4, 1, 2)$ is the instance of *SelectiveCommunication* involving three agents (α_1, α_2 and α_3), four rooms (ρ_1, ρ_2, ρ_3 and ρ_4), one box (β_1) and two possible colors for the box (e.g., red and green). Figure 1 depicts a *possible state* of the environment in this domain.

A possible goal $g = g_1 \wedge g_2 \wedge g_3$ in this domain is given by the conjunction of the following sub-goals:

- g_1 : α_1 and α_2 know the color of β_1
- g_2 : α_1 knows that α_2 knows the color of β_1
- g_3 : α_1 knows that α_3 does not know the color of β_1




Figure 1: A depiction of a possible state in $SC(3,4,1,2)$, where a red box β_1 is in room ρ_2 .

That is, g requires α_1 to learn the color of β_1 and privately communicate this information to α_2 . I.e., the goal is epistemic; it requires α_1 to achieve *first-order knowledge* about the environment (g_1) as well as *higher-order knowledge* about what others know (g_2 and g_3).

The nature of a *plan* for achieving g depends on the *initial state* as well as the assumptions made about the planning problem. For illustrative purposes, we consider a simple problem. Suppose that only α_1 can act and that the initial state s_0 satisfies the following conditions:

- c_1 : each agent knows the location of all agents and the box β_1 .
- c_2 : no agent knows the color of β_1 (which is in fact red).
- c_3 : conditions c_1 and c_2 are common knowledge among the agents.

In this case, α_1 can easily reach a state satisfying g from s_0 . The following plan achieves g : α_1 moves to ρ_2 , α_1 senses the color of β_1 , α_1 announces the color of β_1 . Of course, more initial uncertainty, or allowing other agents to act (sequentially or in parallel), results in more complex tasks. Such tasks can be defined with the formalism presented in this paper; however, for a first take on the formalism, this toy problem will be considered.

For additional examples, in [54] we use the framework to model social networks with epistemic dynamics.

3 Language

As term-modal logical languages include first-order aspects, they are parameterized by a *signature* specifying the *non-logical symbols* and their *type*—i.e., the constants and relation- and function symbols, and the sort and order of arguments (agent or object) they apply to. Also variables are here assigned a type.

Notation 1. For a vector $v = (x_1, \dots, x_n)$, let $len(v)$ denote its length, let v_i denote its i 'th element, i.e., $v_i := x_i$, and let $v|_k$ denote its restriction to its prefix of length k , i.e., $v|_k := (x_1, \dots, x_k)$.

Definition 1. A *signature* is a tuple $\Sigma = (V, C, R, F, \tau)$ with V a countably infinite set of variables, and C , R and F countable sets of respectively constants, relation symbols and function symbols with the one requirement that $\{=\} \subseteq R$. Finally, τ is a type assignment map that satisfies

1. For $x \in V$, $\tau(x) \in \{\text{agt}, \text{obj}\}$ such that both $V \cap \tau^{-1}(\text{agt})$ and $V \cap \tau^{-1}(\text{obj})$ are countably infinite.
2. For $c \in C$, $\tau(c) \in \{\text{agt}, \text{obj}\}$.

3. For $r \in R$,
 - (a) for some $n \in \mathbb{N}$, $\tau(r) \in \{\text{agt}, \text{obj}, \text{agt_or_obj}\}^n$, and
 - (b) for $= \in R$, $\tau(=) = (\text{agt_or_obj}, \text{agt_or_obj})$.
4. For $f \in F$,
 - (a) for some $n \in \mathbb{N}$, $\tau(f) \in \{\text{agt}, \text{obj}, \text{agt_or_obj}\}^n \times \{\text{agt}, \text{obj}\}$, and
 - (b) if $\tau(f)|_n = (\tau(t_1), \dots, \tau(t_n))$, then $\tau(f(t_1, \dots, t_n)) = \tau(f)_{n+1}$.

Identity is treated as a relation symbol; for it, infix notation is used with $= (a, b)$ written $a = b$.

Example 1 (Signature for SelectiveCommunication). The following signature $\Sigma = (V, C, R, F, \tau)$ can be used to specify the $\text{SC}(n, m, k, \ell)$ domain introduced in Section 2:

- Variables $V = \{x^*, x, y, z, x_1, x_2, x_3, \dots\}$.
- Constants $C = \text{Agents}_{\text{con}} \cup \text{Rooms}_{\text{con}} \cup \text{Boxes}_{\text{con}} \cup \text{Colors}_{\text{con}}$, where $\text{Agents}_{\text{con}} = \{a_1, \dots, a_n\}$, $\text{Rooms}_{\text{con}} = \{r_1, \dots, r_m\}$, $\text{Boxes}_{\text{con}} = \{b_1, \dots, b_k\}$ and $\text{Colors}_{\text{con}} = \{c_1, \dots, c_\ell\}$.
- Relation symbols $R = \{\text{In}, \text{Color}, \text{Adj}, =\}$ where $\text{In}(x, y)$ states that agent or box x is in room y , $\text{Color}(x, y)$ states that box x has color y and $\text{Adj}(x, y)$ states that room x is adjacent to room y .
- Function symbols $F = \emptyset$.
- Type assignment τ with constant types $\tau(x) = \text{agt}$ for $x \in \text{Agents}_{\text{con}}$, $\tau(x) = \text{obj}$ for $x \in \text{Rooms}_{\text{con}} \cup \text{Boxes}_{\text{con}} \cup \text{Colors}_{\text{con}}$, relation types $\tau(\text{In}) = (\text{agt_or_obj}, \text{obj})$, and $\tau(\text{Color}) = \tau(\text{Adj}) = (\text{obj}, \text{obj})$.

Definition 2. The set of *terms* T of a signature $\Sigma = (V, C, R, F, \tau)$ is given by the grammar

$$t ::= x \mid c \mid f(t_1, \dots, t_n)$$

for $x \in V$, $c \in C$ and $f \in F$, provided that $t_1, \dots, t_n \in T$ and $\tau(f)|_{\text{len}(\tau(f))-1} = (\tau(t_1), \dots, \tau(t_n))$.

A term is *ground* if it does not contain any variables; it is *free* if all its terms are (i) variables or (ii) function symbols all whose arguments are free terms.

By the definitions of type assignments and terms, it is the case that for all $t \in T$, $\tau(t) \in \{\text{agt}, \text{obj}\}$. This allows for a uniform definition of formulas in term-modal languages:

Definition 3. Let $\Sigma = (V, C, R, F, \tau)$ be a signature. Let $t_1, \dots, t_n \in T$ and $r \in R$ with $\tau(r) = (\tau(t_1), \dots, \tau(t_n))$, let $\dagger \in T$ with $\tau(\dagger) = \text{agt}$, and let $x \in V$. The *language* \mathcal{L} is then given by the grammar

$$\varphi ::= r(t_1, \dots, t_n) \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_{\dagger}\varphi \mid \forall x\varphi$$

An *atom* is a formula obtained by the first clause. An atom is *ground* if all its terms are ground; it is free if all its terms are *free*. Denote by $\text{GroundAtoms}(\mathcal{L})$ and $\text{FreeAtoms}(\mathcal{L})$ the set of all ground and free atoms in \mathcal{L} , respectively.

Throughout, the standard Boolean connectives as well as \top , \perp and \exists are used as meta-linguistic abbreviations as usual. We abbreviate inequality expressions of the form $\neg(t_1 = t_2)$ by $(t_1 \neq t_2)$. *Free* and *bound* variables may be defined recursively as usual with the free variables of $K_t\varphi$ the free variables of φ plus the variables in t . A formula is a *sentence* if it has no free variables. With $\varphi \in \mathcal{L}$, $t \in \mathbb{T}$, $x \in \mathbb{V}$, $\tau(x) = \tau(t)$ and no bound variables of φ occurring in t , the result of replacing all occurrences of x in φ with t is denoted $\varphi(x \mapsto t)$.

Remark 1. $K_t\varphi$ is read as “agent t knows that φ ”. Epistemic expressions are only well-defined when t is an agent term. The language \mathcal{L} neither enforces nor requires a fixed-size agent set, in contrast with standard epistemic languages, where the set of operators is given by reference to some index set. Fixed-size agent sets are discussed throughout.

4 State Representation

In planning frameworks based on epistemic logic, states are often represented using *possible-worlds models*, tracing back to the work of Hintikka [40] and Kripke [50]. The standard epistemic interpretation of such models—employed here in all examples—is one of *indistinguishability*, as follows. A model contains a set of worlds, each representing a physical state of affair. For each agent, a model contains a binary relation on the set of worlds. Under the indistinguishability interpretation, this relation is taken to be an equivalence relation. If two worlds are related for agent α , then α cannot distinguish them given her current information. I.e., they are *informationally indistinguishable* for α . Hence, when α in fact is in some world w , she cannot tell which of the worlds related to w she is in fact in. The set of worlds indistinguishable from w for agent α is therefore sometimes referred to as agent α 's *range of uncertainty* (at w). The term *information cell* is used to cover the same, and a world in α 's range of uncertainty is said to be considered *possible* by α (at w). An agent's range of uncertainty determines its knowledge: an agent knows φ in world w if φ is true in all the worlds in the agent's range of uncertainty at w . For instance, if the agent has no information about two blocks A and B , and therefore cannot tell whether one of them is stacked on the other or not, she will consider at least three worlds possible: one in which A is indeed stacked on B , one in which it is not, and one in which B is stacked on A . Possible-worlds models represent also all levels of *higher-order knowledge*. E.g. agent α knows that agent β knows φ if α does not consider it possible that β considers it possible that φ is false.

A possible-worlds model is formally defined as a structure in general called a *Kripke model*. Kripke models need not enforce any properties on the agents' relations. Our results hold for the general case, with equivalence relations a special case. Under the indistinguishability interpretation, Kripke models are often called *epistemic models* or *epistemic states*. For a thorough explanation of the components of a Kripke model, we refer the reader to [8, 28]. When the context makes it clear, such a structure may simply be called a *model*. A model consists of a *frame* and an *interpretation*. Two things differentiate the frame used here from the standard, propositional version. First, a frame here contains a constant domain of elements existing in each world. Working with distinct agents and objects, the domain is a disjoint union of two sets, the agent domain and the object domain. Second, the accessibility relations over worlds are directly associated with elements in the agent domain. The agent domain thereby makes reference to an index set—as used in non-term-modal logical frames—redundant. The definition of a frame is thereby self-contained.

Definition 4. A *frame* F is a triple $F = (D, W, R)$ where

1. $D := D_{\text{agt_or_obj}} := D_{\text{agt}} \dot{\cup} D_{\text{obj}}$, called the *domain*, is the disjoint union of the non-empty sets D_{agt} and D_{obj} , called the *agent domain* and the *object domain*, respectively.
2. W is a non-empty set of *worlds*.
3. R is a map associating to each agent $i \in D_{\text{agt}}$ a binary *accessibility relation* on W . I.e., $R : D_{\text{agt}} \longrightarrow \mathcal{P}(W \times W)$.

Write R_i for $R(i)$, write wR_iw' for $(w, w') \in R_i$ and write $R_i(w)$ for $\{w' \in W : wR_iw'\}$. If $|D_{\text{agt}}| = n$ and $|D_{\text{obj}}| = m$, ($n, m \in \mathbb{N}$), say F is of *size* (n, m) . Denote by F the class of all frames and by $F_{(n,m)}$ the class of all frames of size (n, m) .

For propositional modal logic, a frame is augmented by a valuation assigning an extension of worlds to each propositional symbol. In the first-order and term-modal cases, each non-logical symbol is assigned an extension in the domain. Here, this extension is assigned *world-relatively* for both relation symbols, function symbols and constants. In particular the last is note-worthy: the constants are thereby *non-rigid*—they may refer to different objects (and agents!) in different worlds. The non-rigidity of constants allows for uncertainty about identity cf. the example of Section 6.2 and play an important role concerning the validity of frame-property characterizing axioms, cf. Section 8.1.4.

Constants that do not vary with worlds—so-called *rigid* constants—often come in handy when referring to agents. A rigid constant provides a syntactic name for a semantic agent. Rigid constants are a special case: a constant c may be forced rigid by assuming its interpretation $I(c, w)$ to be constant over all worlds, i.e., by $I(c, w) = I(c, w')$ for all $w, w' \in W$.

Definition 5. Let a signature $\Sigma = (\mathcal{V}, \mathcal{C}, \mathcal{R}, \mathcal{F}, \tau)$ and a frame $F = (D, W, R)$ be given. An *interpretation* of Σ over F is a map I satisfying for each $w \in W$:

1. $I(=, w)$ is the set $\{(d, d) : d \in D\}$.
2. For $c \in \mathcal{C}$, $I(c, w) \in D_{\tau(c)}$.
3. For $r \in \mathcal{R}$, $I(r, w) \subseteq \prod_{i=1}^{\text{len}(\tau(r))} D_{\tau_i(r)}$.
4. For $f \in \mathcal{F}$, $I(f, w) \subseteq \prod_{i=1}^{\text{len}(\tau(f))} D_{\tau_i(f)}$ such that $I(f, w)$ is a (possibly partial) map: i.e., if $(d_1, \dots, d_k, d_{\text{len}(\tau(f))}) \in I(f, w)$, then $d_{\text{len}(\tau(f))} = d'_{\text{len}(\tau(f))}$.

With $F = (D, W, R)$ a frame and I an interpretation of Σ over F , the tuple $M = (D, W, R, I)$ is a *model*. Both $w \in F$ and $w \in M$ states that $w \in W$. When $w \in M$, the pair (M, w) is a *pointed model* with w called the *designated world*.

Satisfaction for all formulas without variable occurrences may be defined over pointed models. To specify satisfaction for the full language, variables must also be assigned extension. Letting variable valuations be world independent—or *rigid*—trans-world identification of objects and agents may be made using suitable bound variables, cf. e.g. the *de re* knowledge in Example 2.

Definition 6. Let a signature $\Sigma = (\mathcal{V}, \mathcal{C}, \mathcal{R}, \mathcal{F}, \tau)$ and a frame $F = (D, W, R)$ be given. A *valuation* of Σ over F is a map $v : \mathcal{V} \longrightarrow D$ such that $v(x) \in D_{\tau(x)}$. An *x -variant* v' of v is a valuation such that $v'(y) = v(y)$ for all $y \in \mathcal{V}, y \neq x$.

Jointly, an interpretation and a valuation assigns an extension to every term t of Σ relative to every world of a frame. The following involved, but uniform, notation will be used throughout to denote the extension of terms:

Definition 7. Let a signature $\Sigma = (\mathcal{V}, \mathcal{C}, \mathcal{R}, \mathcal{F}, \mathfrak{t})$, a model $M = (D, W, R, I)$ and a valuation v be given. The *extension* of the term $t \in \mathcal{T}$ in M under v is

$$\llbracket t \rrbracket_w^{I,v} = \begin{cases} v(t) & \text{if } t \in \mathcal{V} \\ I(t, w) & \text{if } t \in \mathcal{C} \\ d \text{ with } (d_1, \dots, d_n, d) \in I(f, w) & \text{if } t = f(t_1, \dots, t_n) \end{cases}$$

For exactly the terms t with $\mathfrak{t}(t) = \text{agt}$, $R_{\llbracket t \rrbracket_w^{I,v}}$ is then an accessibility relation in M . The extension of terms thus play a key role in the satisfaction of modal formulas:

Definition 8. Let $\Sigma = (\mathcal{V}, \mathcal{C}, \mathcal{R}, \mathcal{F}, \mathfrak{t})$, $M = (D, W, R, I)$ and v be given. The *satisfaction* of formulas of \mathcal{L} is given recursively by

- $M, w \models_v r(t_1, \dots, t_n)$ iff $(\llbracket t_1 \rrbracket_w^{I,v}, \dots, \llbracket t_n \rrbracket_w^{I,v}) \in I(r, w)$ for all $r \in \mathcal{R}$, including $=$.
- $M, w \models_v \neg \varphi$ iff not $M, w \models_v \varphi$.
- $M, w \models_v \varphi \wedge \psi$ iff $M, w \models_v \varphi$ and $M, w \models_v \psi$.
- $M, w \models_v \forall x \varphi$ iff $M, w \models_u \varphi$ for every x -variant u of v .
- $M, w \models_v K_t \varphi$ iff $M, w' \models_v \varphi$ for all w' such that $(w, w') \in R_{\llbracket t \rrbracket_w^{I,v}}$.

Example 2 (Epistemic model for $\text{SC}(3, 4, 1, 2)$). Figure 2 depicts an epistemic model $M_0 = (D, W, R, I)$ for the initial state s_0 described in Section 2.




Figure 2: (M_0, w_{red}) , a pointed epistemic model for the initial state s_0 described in Section 2. The agents are uncertain about the color of β_1 , which may be red (w_{red}) or green (w_{green}). This is captured by the edge linking w_{red} and w_{green} . Reflexive edges are not drawn. The name of the actual world, w_{red} , is marked with boldface letters.

Formally, the model $M_0 = (D, W, R, I)$ has

- $D = D_{\text{agt}} \dot{\cup} D_{\text{obj}}$, with $D_{\text{agt}} = \{\alpha_1, \alpha_2, \alpha_3\}$ and $D_{\text{obj}} = \text{Rooms}_{\text{obj}} \cup \text{Boxes}_{\text{obj}} \cup \text{Colors}_{\text{obj}}$, where $\text{Rooms}_{\text{obj}} = \{\rho_1, \rho_2, \rho_3, \rho_4\}$, $\text{Boxes}_{\text{obj}} = \{\beta_1\}$ and $\text{Colors}_{\text{obj}} = \{\text{Red}, \text{Green}\}$.
- $W = \{w_{red}, w_{green}\}$.
- $R(\alpha_i) = W \times W$, for $i \in \{1, 2, 3\}$.

- The interpretation of all constants is the same in w_{red} and w_{green} , i.e., all constants are rigid: $I(a_i, u) = \alpha_i$, $I(r_i, u) = \rho_i$, $I(b_i, u) = \beta_i$, $I(green, u) = Green$ and $I(red, u) = Red$ for all $u \in W$.
- The interpretation of the predicates is as follows: $I(In, u) = \{(\alpha_1, \rho_1), (\alpha_2, \rho_3), (\alpha_3, \rho_4), (\beta_1, \rho_2)\}$, for all $u \in W$, $I(Color, w_{red}) = \{(\beta_1, red)\}$, $I(Color, w_{green}) = \{(\beta_1, green)\}$. The interpretation of the Adj predicate is as expected.

Following the semantics from Definition 8, it can be seen that $M_0, w_{red} \models_v \forall x(K_x In(b_1, r_2))$, i.e., every agent knows the location of box β_1 . Similarly, every agent knows that all agents know this, since $M_0, w_{red} \models_v \forall y \forall x(K_y K_x In(b_1, r_2))$. Moreover, the agents know that the box has a color, but not what color it is. They thus have what is called *de dicto* knowledge of the coloring of the box, but not *de re* knowledge. Agent α_3 's de dicto knowledge is captured by $M_0, w_{red} \models_v K_{\alpha_3} \exists x Color(b_1, x)$, while its lack of de re knowledge is captured by $M_0, w_{red} \models_v \neg \exists x K_{\alpha_3} Color(b_1, x)$. Finally, agent α_3 has de re knowledge of the box, or as Hintikka [40] puts it, α_3 knows what the box is, captured by $M_0, w_{red} \models_v \exists x K_{\alpha_3}(x = b_1)$.

5 Action Representation

In automated planning, a distinction is often drawn between *action schemas*, which describe classes of actions in a general way, and *ground actions*, which represent a specific action with a fixed set of agents and objects [36, 80]. Action schemas use so-called *action parameters* or *variables*, which are instantiated into constants to define an action. For example, a schema may be used to represent all actions of the form 'agent x tells y that object z has color u ', where x, y, z and u are variables standing for agents and objects. A corresponding ground action is obtained by replacing all free variables by names referring to specific agents and objects. For example, a schema instance could be 'ann tells bob that box1 has color red'.

In DEL, the descriptions of concrete actions are called *action models*. That is, DEL action models correspond to a ground action in classical planning. Following the DEL naming conventions, models of concrete actions will be called action models, whereas variabilized models in the spirit of PDDL will be called action schemas. Action models and action schemas, as well as a suitable notion of schema instantiation relating the two, are introduced next.

5.1 Action Models

Formally and intuitively, action models are closely related to Kripke models. Where Kripke models contains worlds and relations, action models instead contain *events* and relations. Under the standard epistemic interpretation, the relations again represent indistinguishability and are again assumed to be equivalence relations. Again, this is a special case of the models introduced here.

The DEL-style action models we add to the term-modal logic setting include *preconditions* ([6, 7]), *postconditions* ([11, 15, 27]) as well as *edge-conditions* similarly to [13]. Preconditions specify when an event is executable (e.g., a precondition of opening the door is that it is closed.) Postconditions describe the physical effects of events (e.g., the door is open after the event). Edge-conditions are used to represent how an agent's observation of an action depends on the agent's circumstances. For example, the way in which an agent α_i observes an action performed by agent α_j may depend on α_i 's proximity to α_j or to the objects affected by the action (e.g., to α_i , the events of opening and closing the door are distinguishable if α_i can see or hear the door, but else not). Edge-conditions provide a

general way to describe actions whose observability is context-dependent. The epistemic effects of an action model is encoded by the *product update* operation by which action models are applied to Kripke models (defined below).

In more detail, the components of term-modal action models (Def. 9 below) play the following roles. E represents the set of events that might occur as the action is executed. Q is a map that assigns to each edge $(e, e') \in E \times E$ an edge-condition: a formula with a single free variable x^* . Given a model M describing the situation in which the action is applied, an agent α cannot distinguish e from e' iff the edge-condition from e to e' is true in M when the free variable x^* is mapped to α . Intuitively, if the situation described by the edge-condition is true for α , the way α observes the action does not allow her to tell whether e or e' is taking place. The precondition restricts the applicability of an event e to those states satisfying the precondition formula $\text{pre}(e)$. Precondition formulas contain no free variables to ensure that their effects are conditional only on the model, but not the variable valuation. The postcondition $\text{post}(e)$ describes the physical changes induced by the event. If both $\text{pre}(e)$ and $\text{post}(e)(r(t_1, \dots, t_n))$ are true in a state s of a model M , then the event e occurs, and after its occurrence, $r(t_1, \dots, t_n)$ is true in the updated version of s . That is, $r(t_1, \dots, t_n)$ is a *conditional effect* of event e with condition $\text{post}(e)(r(t_1, \dots, t_n))$.

The language used to state pre- and postconditions in action models is an extension of \mathcal{L} , denoted \mathcal{L}_{AM} , to be introduced in Section 7. This extended language has formulas of the form $[A, e]\varphi$, which are interpreted as: ‘after event e of action A , φ holds’. This type of formula makes it possible to mention other actions in the pre- and postconditions of actions, i.e., to express syntactically some dependencies or interactions between actions. However, the action model construction does not require or depend on the use of \mathcal{L}_{AM} rather \mathcal{L} , so the reader can safely ignore the details of \mathcal{L}_{AM} for now.

Definition 9. An *action model* A is a tuple $A = (E, Q, \text{pre}, \text{post})$ where

1. E is a non-empty, finite set of possible *events*.
2. $Q : (E \times E) \rightarrow \mathcal{L}_{AM}$, where for each pair (e, e') the formula $Q(e, e')$ has exactly one free variable x^* .
3. $\text{pre} : E \rightarrow \mathcal{L}_{AM}$ is a map that assigns to each event $e \in E$ a precondition formula with no free variables.
4. $\text{post} : E \rightarrow (\text{GroundAtoms}(\mathcal{L}) \rightarrow \mathcal{L}_{AM})$ is a map that assigns to each event $e \in E$ a postcondition for each ground atom.

It is required that $\text{post}(e)(= (t, t)) = \top$ for each event e , to preserve the meaning of equality. A pair (A, e) consisting of the action and an event from E is called a *pointed action*.

Notation 2. Let $A = (E, Q, \text{pre}, \text{post})$ be an action model. We denote by $\text{dom}(\text{post}(e))$ the set of atoms for which $\text{post}(e)(r(t_1, \dots, t_k)) \neq r(t_1, \dots, t_k)$. We denote any $\text{post}(e)$ that maps every atom to itself by *id* (the identity function). When convenient, we add the superscript “ A ” to the components of A , so that $A = (E^A, Q^A, \text{pre}^A, \text{post}^A)$.

To ensure that postcondition functions are finite objects, each $\text{post}(e)$ is often required to be only finitely different from the identity function. That is, $\text{dom}(\text{post}(e))$ is required to be finite. This allows for a finite encoding of postconditions, as only pairs with $\text{post}(e)(\varphi) \neq \varphi$ need to be stored in memory. Especially in planning, it should be possible to write down a sequence of symbols that completely specifies any given action model in the language. For the sake of generality, we do not impose this

restriction in the definition of an action model. But for all practical purposes, this standard restriction will be needed.

Notation 3. Let $A = (E, Q, \text{pre}, \text{post})$ be an action model. When A is illustrated as a labelled graph, for each node $e \in E$, we write the precondition and postconditions for e as a pair $\langle \text{pre}(e); \text{post}(e)(\psi_1) = \varphi_1 \wedge \dots \wedge \text{post}(e)(\psi_n) = \varphi_n \rangle$. We write postconditions such as $\text{post}(e)(\varphi) = \top \wedge \text{post}(e)(\psi) = \perp$ using the notation $\varphi \wedge \neg\psi$ (indicating that the action makes φ true and ψ false unconditionally). In graphs, we omit the postconditions for atoms φ with $\text{post}(e)(\varphi) = \varphi$.

We do not include the edge-conditions for reflexive loops in illustrations, but always assume that for all $e \in E$, $Q(e, e) = (x^* = x^*)$, to the effect that all agents retain reflexive relations following updates. When two events e, e' are connected by a line without arrowheads labeled by a single formula φ , this means that $Q(e, e') = Q(e', e) = \varphi$, retaining symmetry.

Example 3 (Action models for $\text{SC}(3, 4, 1, 2)$). Figures 3, 4 and 5 depict graphically the action models for the three actions in the plan from in Section 2, i.e., the following movement, sensing and announcement actions: α_1 moves to ρ_2 , α_1 senses the color of β_1 , α_1 announces the color of β_1 .




Figure 3: $\text{Move}(a_1, r_1, r_2)$, the action model for α_1 moving from ρ_1 to ρ_2 . Event e_m describes what is actually taking place (in the drawing for the model, the actual event is marked with a double circle). The precondition formula says that α_1 is in ρ_1 and that ρ_2 is next to ρ_1 . The action changes α_1 's location to ρ_2 , as captured by the postcondition. The event e'_m describes the situation in which nothing happens. This is how the action looks to any agent that is neither in the room α_1 is currently in, nor in the room the agent is moving to. The edge-condition linking the two events captures this observability constraint.




Figure 4: $\text{SenseCol}(a_1, \text{red}, b_1, r_2)$, the action model for α_1 sensing in room ρ_2 whether box β_1 is red or not. Event e_s describes what is actually taking place, i.e., α_1 seeing that the box is red. The action is a purely epistemic action, i.e., it does not change the physical state of the environment, and therefore the postcondition $\text{post}(e_s)$ is id . e'_s represents the event in which α_1 sees that the box is not red, while e''_s represents the event in which nothing happens. The agents that are not in ρ_2 cannot observe what α_1 is doing. More precisely, they cannot distinguish between α_1 seeing that the box is red, α_1 seeing that it is not red, and α_1 doing nothing. This is captured by the edge-conditions.




Figure 5: $\text{Announce}(r_1, \text{red}, b_1, r_2)$, the action model for α_1 announcing that β_1 is red while in room ρ_2 . Event e_a describes what is actually taking place. The precondition formula $\text{pre}(e_a)$ says that α_1 is in ρ_2 and that a_1 knows that the color of β_1 is red. Event e'_a describes the event in which nothing occurs. This is what the announcement looks like to any agent that cannot hear the announcement. An agent α_i cannot hear α_1 's announcement if α_i is neither in α_1 's room nor in a room that is adjacent to it. This is captured by the (identical) edge-conditions $Q(e_a, e'_a)$ and $Q(e'_a, e_a)$.

5.2 Product Update

Having defined epistemic models and action models, we introduce an operation that computes the epistemic model M' reached by applying action A in model M . The operation is a first-order variant of *product update* [7]. Under the indistinguishability interpretation, the core epistemic intuition is that to tell two worlds apart after an update, either the agent could tell them apart beforehand, or it could tell them apart by something happening in one, but not the other. In slightly more detail: Assume that *after* an update, a model contains worlds (w, e) and (w', e') , representing that event e occurred in world w , and e' occurred in w' . Then (w, e) is indistinguishable from (w', e') for agent α iff α found both w and w' indistinguishable and events e and e' indistinguishable. Formally, (term-modal) product update is defined below. An explanatory remark follows the definition.

Definition 10. Let $M = (D, W, R, I)$ and $A = (E, Q, \text{pre}, \text{post})$ be given. The *product update* of M and A yields the epistemic model $M \otimes A = (D', W', R', I')$ where

1. $D' = D$
2. $W' = \{(w, e) \in W \times E : (M, w) \models_v \text{pre}(e)\}$,
3. For each $i \in D_{\text{agt}}$, $(w, e)R'_i(w', e')$ iff $wR_i w'$ and $M, w \models_{v[x^* \mapsto i]} Q(e, e')$,
4. $I'(c, (w, e)) = I(c, w)$ for all $c \in \mathcal{C}$, $I'(f, (w, e)) = I(f, w)$ for all $f \in \mathcal{F}$, and $I'(r, (w, e)) = (I(r, w) \cup r^+(w)) \setminus r^-(w)$, where:

$$r^+(w) := \{(\llbracket t_1 \rrbracket_w^{I, v}, \dots, \llbracket t_k \rrbracket_w^{I, v}) : (M, w) \models_v \text{post}(e)(r(t_1, \dots, t_k))\}$$

$$r^-(w) := \{(\llbracket t_1 \rrbracket_w^{I, v}, \dots, \llbracket t_k \rrbracket_w^{I, v}) : (M, w) \not\models_v \text{post}(e)(r(t_1, \dots, t_k))\}$$

If $(M, w) \models_v \text{pre}(e)$, (A, e) is **applicable** to (M, w) . If (A, e) is applicable to (M, w) , the product update of the two yields the pointed epistemic model $(M \otimes A, (w, e))$. Else it is undefined.

Remark 2. The components of the updated model are as follows. The domain of the updated model D' is unchanged, since action models change the state of agents and objects, but do not introduce or remove them. A state (w, e) is in the updated set of states W' if, and only if, e is applicable in w , i.e., if (M, w) satisfies the precondition $\text{pre}(e)$. As $\text{pre}(e)$ has no free variables by construction, the set of worlds W' is independent of the assignment v . The state (w, e) represents the state reached by taking event e in state w . Agent α cannot distinguish (w, e) from (w', e') if (1) α cannot distinguish w from w' , which is the case if $wR_\alpha w'$; and (2) α cannot distinguish e from e' given its circumstances in w , which is the case if the edge-condition $Q(e, e')$ is true for agent α at (M, w) when x^* is mapped to α , i.e., when $M, w \models_{v[x^* \mapsto i]} Q(e, e')$. Since actions do not change the denotation of ground terms, I' agrees with I in this respect. The extension of relations is changed according to event postconditions. If the condition $\text{post}(e)(r(t_1, \dots, t_k))$ is true at (M, w) , then the tuple $(\llbracket t_1 \rrbracket_w^{I, v}, \dots, \llbracket t_k \rrbracket_w^{I, v})$ is added to the extension of r at (w, e) , and it is removed otherwise.

Example 4 (Product updates for $\text{SC}(3, 4, 1, 2)$). Starting from the initial epistemic model (M_0, w_{red}) from Example 2, we model the effects of applying the actions in α_1 's plan. First, α_1 moves right. This action yields the new pointed model $(M_0 \otimes \text{Move}(a_1, r_1, r_2), (w_{red}, e_m))$, depicted in Figure 6.

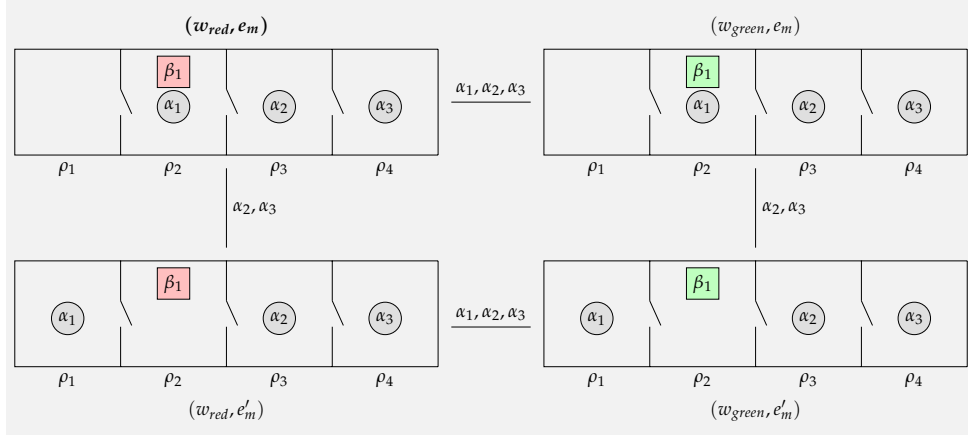


Figure 6: The pointed model $(M_0 \otimes \text{Move}(a_1, r_1, r_2), (w_{red}, e_m))$, representing the state after α_1 moves into room ρ_2 . Edges in the reflexive-transitive closure of the indistinguishability relations are omitted. At this point, α_2 and α_3 are uncertain about the location of α_1 . More precisely, they cannot tell whether α_1 stayed in room ρ_1 or moved to ρ_2 .

The second step is sensing the color of β_1 . This action yields the model $(M_0 \otimes \text{Move}(a_1, r_1, r_2) \otimes \text{SenseCol}(a_1, red, b_1, r_2), (w_{red}, e_m, e_s))$, depicted in Figure 7.




Figure 7: The pointed model $(M_0 \otimes \text{Move}(a_1, r_1, r_2) \otimes \text{SenseCol}(a_1, \text{red}, b_1, r_2), (w_{\text{red}}, e_m, e_s))$, representing the state after α_1 senses that β_1 is red while in room ρ_2 . Edges in the reflexive-transitive closure of the indistinguishability relations are omitted. Note that, in the actual world, $(w_{\text{red}}, e_m, e_s)$, α_1 does not face any uncertainty. In particular, $K_{\alpha_1} \text{Color}(b_1, \text{red})$ is true at $(w_{\text{red}}, e_m, e_s)$. On the other hand, α_2 and α_3 have not observed any of α_1 's actions. That is, α_1 may or may not have moved, and it may or may not have sensed whether β_1 is red. As a result, it holds at $(w_{\text{red}}, e_m, e_s)$ that $\forall x(x \neq a_1 \rightarrow \neg K_x \text{In}(a_1, r_2) \wedge \neg K_x \text{Color}(b_1, \text{red}))$.

Finally, α_1 announces in room ρ_2 that the color of β_1 is red. The result of this action is the model $(M_0 \otimes \text{Move}(a_1, r_1, r_2) \otimes \text{SenseCol}(a_1, \text{red}, b_1, r_2) \otimes \text{Announce}(a_1, \text{red}, b_1, r_2), (w_{\text{red}}, e_m, e_s, e_a))$, depicted in Figure 8.

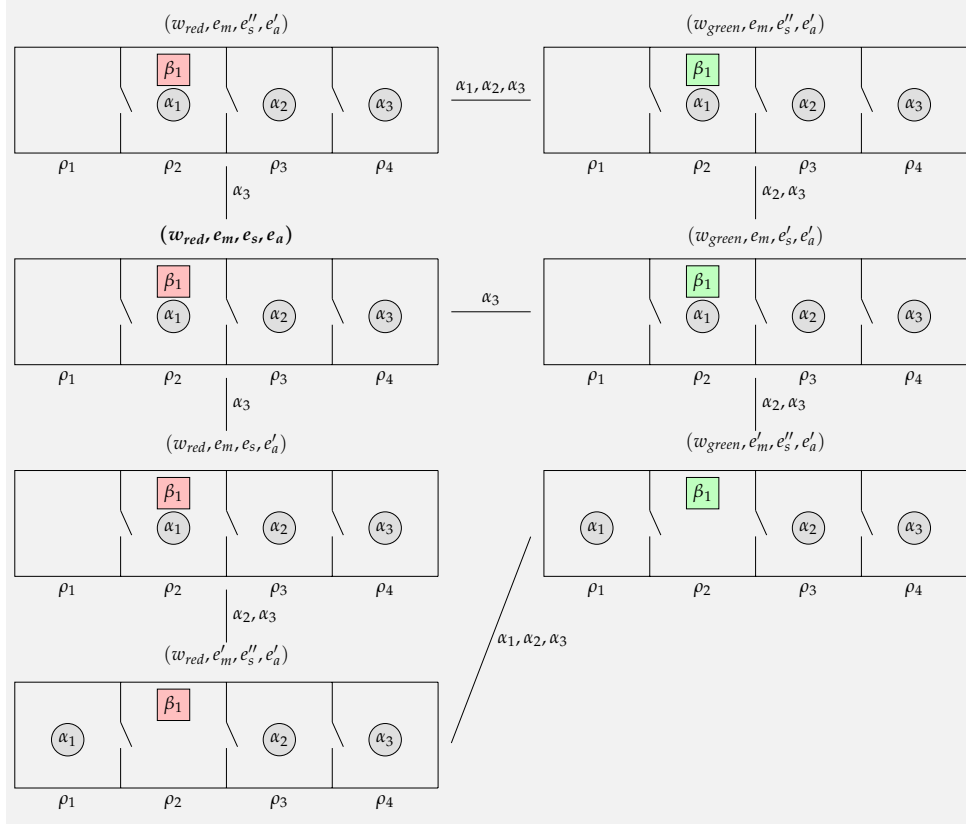


Figure 8: The pointed model representing the state after α_1 announces that β_1 is red while at room ρ_2 , $(M_0 \otimes \text{Move}(a_1, r_1, r_2) \otimes \text{SenseCol}(a_1, \text{red}, b_1, r_2) \otimes \text{Announce}(a_1, \text{red}, b_1, r_2), (w_{\text{red}}, e_m, e_s, e_a))$. Edges in the reflexive-transitive closure of the indistinguishability relations are omitted. Note that, in the actual world, $(w_{\text{red}}, e_m, e_s, e_a)$, α_1 and α_2 do not face any uncertainty; there are no outgoing edges from the actual world for these agents. The goal g as stated in Section 2 holds in the actual world: both α_1 and α_2 know the color of β_1 , α_1 knows that α_2 knows this, and α_1 knows that α_3 does not know this.

5.3 Succinct Representation of Actions via Epistemic Action Schemas

We introduce epistemic action schemas, which represent sets of actions in a general way, as done in common planning formalisms such as PDDL. Schemas use *variables* to describe actions, rather than *constant symbols*. These variables denote arbitrary agents and objects and are used to describe their roles with respect to a type of action, such as the roles of *speaker* and *listener* in an action of type ‘announcement’.

As anticipated in Section 3, a major reason for introducing schemas is that they result in action representations whose size is independent of the number of agents and objects in a domain. For the *SelectiveCommunication* domain $\text{SC}(n, m, k, \ell)$, there are $n \cdot m \cdot k \cdot \ell \cdot 2^{n-1}$ possible announcement actions, since each of the n agents could, in each of the m rooms, announce about each of the k boxes,

that it is of one out of ℓ colors, with one out of the 2^{n-1} subsets of the other agents hearing the announcement. Representing all actions requires $n \cdot m \cdot k \cdot \ell \cdot 2^{n-1}$ standard DEL action models, i.e., one model per action. Variants of standard DEL models such as edge-conditioned models [13] fare substantially better, since the set of hearers is implicitly represented in such models, but $n \cdot m \cdot k \cdot \ell$ models are still required to represent the set of announcements. Other variants of DEL are also more succinct than standard DEL action models, e.g. the symbolic models of [22, 87]. However, all these announcements can be compactly represented with a single *epistemic action schema*, as shown below in Example 5.

Moreover, epistemic schemas open up the possibility of applying well-known techniques such as *least commitment* or *partial order planning* [90] to epistemic problems. These approaches use the notion of a *partially instantiated action*, such as $\text{Move}(B, x, C)$, where x is a variable whose substitution has not yet been chosen. If specifying a binding constraint for x is unnecessary at the current point in the planning process, it is often advantageous to delay this commitment until later, i.e., until other necessary parts of the plan are discovered that further constrain what x should be. Other approaches to lifted planning, such as *hierarchical task networks* (HTNs) ([36], ch. 11) similarly exploit partial substitutions to optimise the search for solutions. For an epistemic version of lifted HTN planning, the epistemic schemas defined here could play the role of primitive tasks.

An epistemic action schema $a(x_1, \dots, x_n)$ is defined using an *action name* a and a *parameter list* (x_1, \dots, x_n) , as done e.g. in PDDL. The parameter list fixes a finite set of agents and objects involved in the execution of the action. Schemas are required to follow a STRIPS-like scope assumption; all variables referenced in the preconditions or postconditions of an action schema must appear in the action's parameter list. Any agent or object unmentioned in the parameter list is assumed to be unrelated to the action's pre- and postconditions.

Definition 11. An *epistemic action schema* is of the form $a(\vec{x}) = (E, Q, \text{pre}, \text{post})$ where

1. a is the *action name* and $\vec{x} \in V^n$ is a finite *parameter list*.
2. E is a non-empty, finite set of *events*.
3. $Q : (E \times E) \rightarrow \mathcal{L}_{AM}$ is an *edge-condition function*, where the formula $Q(e, e')$ has a free variable x^* of type *agt*, and possibly other free variables all in \vec{x} .
4. $\text{pre} : E \rightarrow \mathcal{L}_{AM}$ assigns to each event a *precondition formula* with all free variables in \vec{x} .
5. $\text{post} : E \rightarrow (\text{FreeAtoms}(\mathcal{L}) \rightarrow \mathcal{L}_{AM})$ assigns to each event a *partial postcondition function* such that if y_1, \dots, y_m all occur in \vec{x} , then $\text{post}(e)(r(y_1, \dots, y_m))$ has all free variables from \vec{x} ; else, $\text{post}(e)(r(y_1, \dots, y_m))$ is undefined.

$\text{dom}(\text{post}(e))$ denotes the set of atoms for which $\text{post}(e)(r(t_1, \dots, t_k)) \neq r(t_1, \dots, t_k)$.

The postcondition for each event e is defined as a partial function, with all atoms whose arguments are not a subset of those occurring in \vec{x} left unaffected. Since the parameter list is required to be finite, this yields a finite encoding of postconditions.

Note that the parameter list of a schema may include *agent variables*. Just like any other action parameter, these variables can appear in the preconditions or effects of the schema. This is in line with what occurs in multi-agent extensions of PDDL such as MAPL or MA-PDDL [19, 49]. Such variables are included to enable the schematization of actions also with respect to agents. In this paper, we adopt the epistemic operators from term-modal logic, *indexed by agent variables*, to achieve

schematization with respect to agents when formalizing epistemic planning. In order to be able to express epistemic pre- or postconditions relative to an agent variable x in an action schema, it is necessary to use the variable in the scope of a modal operator. For example, if an action has e.g. a precondition that requires agent x to know $P(y, z)$, we need the term-modal formula $K_x P(y, z)$ to express this constraint in a schematic way.

Example 5 (Action schemas for $SC(n, m, k, \ell)$). Figures 9, 10 and 11 depict graphically the action schemas for the movement, sensing and announcement actions described in the example from Section 2. The schemas have the same structure as the action models from Example 3, but the conditions are now expressed in a general way, via free variables for agents, boxes and colors.




Figure 9: $\text{Move}(x, y, z)$, the action schema for agent x moving from room y to room z .




Figure 10: $\text{SenseCol}(x_1, x_2, x_3, x_4)$, the action schema for x_1 sensing in room x_4 whether box x_3 has color x_2 .




Figure 11: $\text{Announce}(x_1, x_2, x_3, x_4)$, the action schema for x_1 announcing that x_3 has color x_2 while in room x_4 .

As usual in planning, schemas can be instantiated into concrete actions via grounding substitutions. Schema instantiation is defined as follows. Let $a(x_1, \dots, x_n)$ be a schema and $\sigma : \{x_1, \dots, x_n\} \rightarrow \mathcal{C}$ be a **grounding substitution**, i.e., a mapping from variables into constants. For a formula φ , let $\varphi\sigma$ be the result of replacing each occurrence of a free variable y in φ by $\sigma(y)$.

Definition 12. Let $a(x_1, \dots, x_n) = (E, Q, \text{pre}, \text{post})$ be an action schema and let $\sigma : \{x_1, \dots, x_n\} \rightarrow \mathbb{C}$ be a grounding substitution. The *action model induced by* σ is $a(\sigma(x_1), \dots, \sigma(x_n)) = (E', Q', \text{pre}', \text{post}')$ where

1. $E' = E$.
2. for each $e, e' \in E$, $Q'(e, e') = Q(e, e')$
3. for each $e \in E$, $\text{pre}'(e) = \text{pre}(e)\sigma$.
4. for each $e \in E$,

$$\text{post}'(e)(r(t_1, \dots, t_n)\sigma) = \begin{cases} \text{post}(e)(r(t_1, \dots, t_n))\sigma & \text{if } \text{post}(e)(r(t_1, \dots, t_n)) \text{ is defined} \\ r(t_1, \dots, t_n)\sigma & \text{otherwise.} \end{cases}$$

Example 6. The announcement model A_1 in Figure 5 is the ground action of schema S_1 from Figure 11 induced by the substitution $\sigma = \{x_1 \mapsto a_2, x_2 \mapsto r_3, x_3 \mapsto b_1, x_4 \mapsto \text{green}\}$.

6 Problems, Plans and Solutions

This section defines *first-order epistemic planning tasks*. An epistemic planning task consists of an initial state, a set of actions, and a goal to be achieved. To solve an epistemic planning task, one may take either an *external* or an *internal* perspective [3]. The external perspective is the view of the system designer, who knows the precise initial state and actual effect of every action. The internal perspective is the view of an in-system agent, who has uncertainty about the state of the world and therefore uncertainty about the effects of executed actions. In the DEL planning framework, an external planning task is defined as a special case of a classical planning task (as in, e.g., [4, 17]). Following [36], any *classical planning domain* can be described as a state-transition system $T = (\mathcal{S}, \mathcal{A}, \gamma)$ where \mathcal{S} is a finite or recursively enumerable set of finite states, \mathcal{A} is a finite set of actions and $\gamma : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is a partial, computable state-transition function. A *classical planning task* is a triple (T, s_0, S_G) , where T is a state-transition system, $s_0 \in \mathcal{S}$ is the initial state and $S_G \subseteq \mathcal{S}$ is the set of goal states. A *solution* to a classical planning task (T, s_0, S_G) is a plan consisting of a finite sequence of actions a_1, a_2, \dots, a_n such that (1) For all $i \leq n$, $\gamma(\gamma(\dots \gamma(\gamma(s_0, a_1), a_2), \dots, a_{i-1}), a_i)$ is defined, and (2) $\gamma(\gamma(\dots \gamma(\gamma(s_0, a_1), a_2), \dots, a_{n-1}), a_n) \in S_G$. Epistemic planning tasks can be defined as special cases of classical planning tasks, as follows.

Definition 13. Let \mathcal{A} be a finite set of action schemas. A *(first-order) epistemic planning task* based on \mathcal{A} is a triple $P = (s_0, \mathcal{A}, \varphi_G)$ where the initial state s_0 is a finite epistemic state with a finite domain, \mathcal{A} is the set of all ground instances of the schemas in \mathcal{A} , and the goal formula φ_G is a sentence of \mathcal{L} . Any epistemic planning task $(s_0, \mathcal{A}, \varphi_G)$ induces a classical planning task $((\mathcal{S}, \mathcal{A}, \gamma), s_0, S_G)$ given by:

- $\mathcal{S} := \{s_0 \otimes a_1 \otimes \dots \otimes a_n \mid n \in \mathbb{N}, a_i \in \mathcal{A}\}$
- $S_G := \{s \in \mathcal{S} \mid s \models \varphi_G\}$
- $\gamma(s, a) := s \otimes a$ if a is applicable in s , else undefined.

A *solution* to an epistemic planning task is a solution to the induced classical planning task.

All the ingredients in this definition of an external planning task can come from the formalism presented here. Note that the planner-modeler in such a task is not one of the agents in the domain D_{agt} . The planner-modeler has access to the actual states s_i , i.e., to pointed models (M, w) where w is the actual world.

Formalisms of internal epistemic planning based on DEL are often defined from the external planning model, either by adding structure to the models or making small modifications. For example, [2] represents internal perspectives using *information cells*, which are defined from the accessibility relations of an epistemic model. An alternative involves using *multi-pointed models* or adding a set of so-called *designated points* to the epistemic model, with each point describing a world that the agent considers as possible from its internal perspective (see e.g. [15]). A third approach uses a belief state representation of the agent’s internal view as primitive and then defines an epistemic model from it [46]. The approach in [3] offers two different flavors of internal view, both defined on the basis of a standard epistemic model. These various notions of internal perspective, as well as their associated planning tasks, may be upgraded to our framework without major modifications. We believe that the simplest way to add internal perspectives to the present would be to adopt the approach proposed by Bolander and Andersen [16]. Save for the fact that such models are multi-pointed, the core semantics remains the same. We therefore envision no major difficulty in bringing the internal perspective into our formalism; similarly defined multi-pointed structures should suffice. We do not develop the detail here.

6.1 Decidability of the Plan Existence Problem

Having defined first-order epistemic planning tasks, a natural first question is whether the corresponding plan existence problem is decidable. We follow Aucher and Bolander [4] in defining the plan existence problem:

Definition 14. Let $n \in \mathbb{N}$. $\text{PlanEx}(n)$ is the problem: “Given a (first-order) epistemic planning task $P = (s_0, \mathcal{A}, \varphi_G)$ where s_0 is an n -agent epistemic state, does P have a solution?”.

For propositional DEL, the corresponding problem is undecidable in general [16]. This entails that the unrestricted first-order problem is undecidable as well, since first-order epistemic planning extends propositional DEL planning. However, decidable and reasonably expressive fragments of propositional DEL planning have been found, such as single-agent planning and multi-agent planning with non-modal preconditions. In [55], we show that *the corresponding first-order fragments are also decidable*. In that paper, bisimulations for the term-modal models presented here are introduced, and shown to have standard model-theoretic properties. Such bisimulations are key in proving the decidability results, as they allow us to show that the state spaces for certain planning fragments are finitely representable, up to bisimulation. We state the main results here and refer the reader to [55] for details.

Theorem 1 ([55]). $\text{PlanEx}(1)$ (*single-agent planning*) is decidable.

Theorem 2 ([55]). *If all actions have non-modal preconditions, then $\text{PlanEx}(k)$ is decidable, for $k \geq 1$.*

6.2 An Example with a PDDL-like Description

With the above, the dynamic term-modal planning framework of the paper has been introduced. This section contains a second example, using its different components in one place. The example

also serves to illustrate how uncertainty about names may play a role in epistemic planning, and how a term-modal planning domain and an associated planning problem may be described using a ‘PDDL-like syntax’. This description is meant as an indication of how such definitions could be standardized with a PDDL flavor, but no attempt is made at a precise syntax.

6.2.1 The *MachineMalfunction* (MM) Domain

In the $MM(n, m, k)$ domain, there are $n + m$ agents supervising k machines, the agents’ tasked to ensure that machines function correctly at all times, where a choice of n, m and k fixes the universe of the domain. The agents have different roles: n of the agents are monitoring the machines for potential malfunctions, while the remaining m agents are system administrators, that from behind a terminal may solve any malfunction by issuing a reboot command to the affected machine. To reboot a machine, admins need to know its serial number, which a monitor may be uncertain about. Hence, the optimal sequence of actions to fix a malfunction problem will depend on how the knowledge of serial numbers is distributed amongst agents. Finally, company policy dictates epistemic preconditions for the actions: monitoring agents are only allowed to report machine as malfunctioning once they know that it is malfunctioning. To avoid deadlock, the requirement for admins is weaker: they may reboot any machine once they know some machine is malfunctioning.

The remainder of this section concerns an external epistemic planning task $MM_task := (s_0; \mathbf{A}; \varphi_g)$ in $MM(1, 2, 2)$, i.e., with one monitor, two admins, and two machines.

6.2.2 Initial State and Goal

Figure 12 depicts the the pointed epistemic model $s_0 := (M_0, w_0) := ((D, W, R, I), w_0)$, the initial state of MM_task .

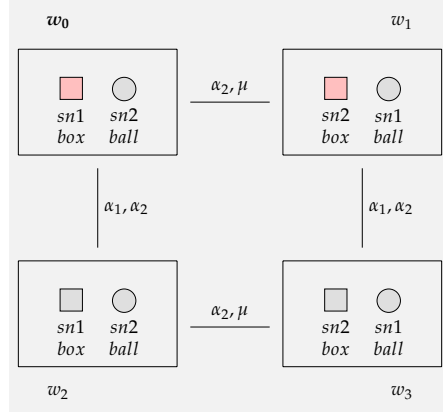


Figure 12: The pointed epistemic model $s_0 = (M_0, w_0)$, the initial state of MM_task. The agent domain D_{agt} comprises three agents, α_1, α_2 and μ , for simplicity only depicted as edge labels. Agents α_1 and α_2 are admins, while μ is a monitor. There are two machines, depicted as a box and a ball; beneath them are the constants that denote them in that world. In w_0 , the box machine is denoted in-system by the serial number $sn1$, but colloquially as *box*, etc. In all worlds, the constants $a1, a2$ and $m1$ denote respectively α_1, α_2 and μ . The red coloring specifies malfunction. The monitor can observe malfunctions, while the admins cannot. Neither the monitor nor the newly employed admin α_2 know the serial numbers, but α_1 does. They all know the colloquial descriptions. The monitor knows that both admins face uncertainty about the malfunction, and that α_1 knows the serial numbers while α_2 does not.

We aim for a simple example. Hence, the administrator/monitor roles are not formally specified, but could straightforwardly be assigned using predicates.

Before stating the available actions, specify the goal of MM_task to be that some agent knows that all machines are not malfunctioning. I.e.,

$$\varphi_g := \exists x K_x \forall y \neg \text{Malfunction}(y)$$

With this goal achieved, the agent knowing that no malfunction occurs can announce this to the remaining agents to achieve that this becomes known to all, but for simplicity, we have omitted this aspect.

The initial state s_0 and the goal φ_g may then be described in a ‘PDDL-like syntax’, as it would appear in a problem file, cf. Figure 13.

```

(define (problem machine-malfunction-p1)
  (:domain machine-malfunction)
  (:universe
    Alpha1 Alpha2 - admin_agent
    Mu           - monitoring_agent
    Box Ball     - machine)
  (:constants
    sn1 sn2 box ball - machine_id
    m1           - monitoring_agent_id
    a1 a2       - admin_agent_id)
  (:init
    (:actual_world w0
      :constant_map ((sn1 Box) (sn2 Ball) (box Box) (ball Ball) (m1 Mu) (a1 Alpha1) (a2 Alpha2))
      :atoms ((malfunction sn1) (malfunction box)))
    (:world w1
      :constant_map ((sn2 Box) (sn1 Ball) (box Box) (ball Ball) (m1 Mu) (a1 Alpha1) (a2 Alpha2))
      :atoms ((malfunction sn2) (malfunction box)))
    (:world w2
      :constant_map ((sn1 Box) (sn2 Ball) (box Box) (ball Ball) (m1 Mu) (a1 Alpha1) (a2 Alpha2))
      :atoms ())
    (:world w3
      :constant_map ((sn2 Box) (sn1 Ball) (box Box) (ball Ball) (m1 Mu) (a1 Alpha1) (a2 Alpha2))
      :atoms ())
    (:edges
      :Alpha1 ((w0 -- w2) (w1 -- w3))
      :Alpha2 (all)
      :Mu ((w0 -- w1) (w2 -- w3)))
    (:goal (exists (?a -agent_id) (knows (?a) (forall (?o - object_id) (not (malfunction ?o)))))))

```

Figure 13: A ‘PDDL-style syntax’ description of the initial state s_0 and the goal φ_0 . A universe of agents and objects is defined using the keyword `:universe`, with constants denoting this domain declared the keyword `:constants`. The `:init` keyword precedes the description of the initial state, which comprises worlds and edges between worlds. Each world is declared with a keyword `:world` and encompasses a `:constant_map` stating what each constant refers to as a list of pairs (constant, entity), as well as a list of true ground `:atoms` (where the closed-world assumption holds). The actual world is defined with the `:actual_world` keyword. The indistinguishability relation for agents is specified with the `:edges` keyword. For each agent, a set of pairs of worlds in the relation is listed, whose reflexive-transitive closure defines the full relation. Finally, the `:goal` keyword declares the goal.

6.2.3 Available Actions as Action Schemas and Domain Definition

To finalize the specification of the external epistemic planning task `MM_task`, the available actions of reporting malfunctions rebooting machines must be defined. To this end, we use the action schemas depicted graphically in Figure 14.




Figure 14: *Left*: $\text{Malfunction}(x, y)$, the action schema for agent x announcing that they know that y is malfunctioning. *Right*: $\text{Reboot}(x, y)$, the action schema for agent x rebooting machine y , an action that x is permitted to do only if x knows that some malfunction is occurring, and which is done privately: other agents than x are uncertain about its execution.

Continuing with the ‘PDDL-like syntax’, Figure 15 describes the *MachineMalfunction* domain. The figure is suggestive of a possible approach to standardizing domain definitions, but again, no formal specification of the syntax is attempted.

```
;; machine-malfunction domain.
(define (domain machine-malfunction)
  (:types admin_agent_id - agent_id
    monitoring_agent_id - agent_id
    serial_number - machine_id
    machine_id
    agent_id)
  (:predicates (malfunction ?o - machine_id))
  (:action MALFUNCTION
    :agent ?s - monitoring_agent_id
    :parameters (?o - machine_id)
    (:actual_event em
      :precondition (knows (?s) (malfunction ?o))
      :postcondition (id))
    (:edge-conditions
      :em -- em (= ?x* ?x*))
  (:action REBOOT
    :agent ?a - admin_agent_id
    :parameters (?n - serial_number)
    (:actual_event er1
      :precondition (knows (?a) (exists (?x - object) malfunction (?x)))
      :postcondition ((malfunction ?n if FALSE))
    (:event er2
      :precondition (TRUE)
      :postcondition (id))
    (:edge-conditions
      :er1 -- er1 (= ?x* ?x*)
      :er2 -- er2 (= ?x* ?x*)
      :er1 -- er2 (not (= ?x* ?a)))
```

Figure 15: A domain definition for the MM domain in a ‘PDDL-style syntax’. Each action schema includes an `: agent` executing it as well as the schema’s `: parameters` list. The possible events comprised in the action and their corresponding edge-conditions are listed next. The actual event is defined with the `: actual_event` keyword, under which pre- and postconditions are listed. Postconditions are given as a list of statements of the form “ground_atom if condition”. A similar `: event` keyword is used for non-actual events. The keyword `edge-conditions` lists, for each pair of events, its edge-condition, via a line of the form `event1–event1(edge-condition)`. Both the `MALFUNCTION` and the `REBOOT` schemas have epistemic preconditions, which can be schematized thanks to the variable-indexed epistemic operators K_x .

6.2.4 Plan and Execution

Given the initial state and actions of the external epistemic planning task MM_task , the goal $\varphi_g = \exists x K_x \forall y \neg \text{Malfunction}(y)$ may be achieved by the monitor agent reporting that the machine colloquially called box is malfunctioning, after which the administrator α_1 may reboot that machine, by knowing its serial number $sn1$. I.e., the plan

$$\pi := \text{Malfunction}(m1, box) \otimes \text{Reboot}(a1, sn1)$$

achieves the goal. Figure 16 depicts the state updates resulting from executing plan π in the state s_0 . In the final state, the goal formula is satisfied, as α_1 knows that no robots is malfunctioning.

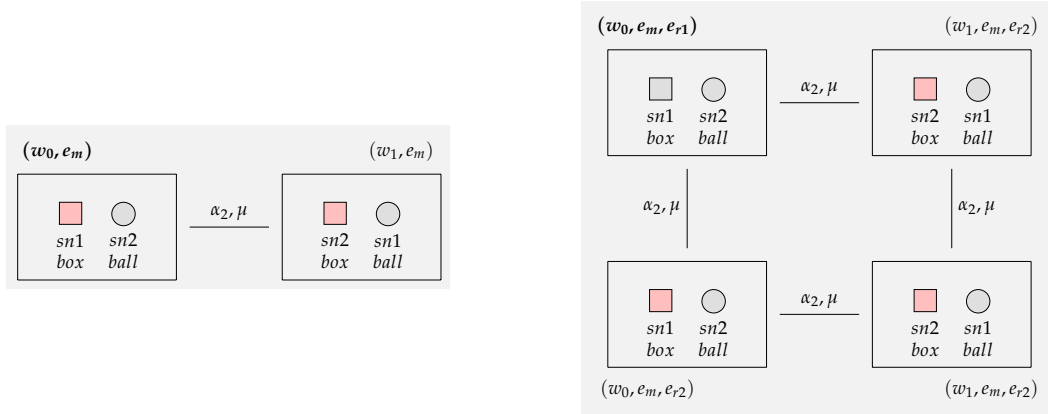


Figure 16: Executing the plan π in s_0 . *Left*: The epistemic state $s_1 := s_0 \otimes \text{Malfunction}(m1, box)$ reached after μ reports that box is malfunctioning. All agents learn that a machine is malfunctioning. However, α_2 and μ are still unsure about the serial number of the malfunctioning robot, but α_1 is not. *Right*: The epistemic state $s_2 := s_1 \otimes \text{Reboot}(a1, box)$ reached after admin agent α_1 reboots $sn1$. Since α_1 has rebooted $sn1$ privately, α_2 and μ still do not know that all machines are functioning. However, the goal is achieved: there is *some* agent that knows that all machines are functioning.

6.2.5 Knowing Who and Alternative Plans

There are other plans than π that would also solve the external epistemic planning task MM_task , but π achieves the goal in the fewest number of steps.

The plan π is kept short by making admin α_1 act and finally witness the existence criteria in the goal φ_g . Admin α_1 is well-suited to this purpose because α_1 does not face uncertainty about the two names box and $sn1$. For the both names, α_1 *knows who* the names refer to, which in turn entails that in s_1 , α_1 knows what machine to reboot.

Any plan where admin α_2 , instead of α_1 , acts, will necessarily be longer, because α_2 does not know is uncertain about what serial number belong to which machine, and must therefore reboot *both* machines. More specifically, then given s_0 , any successful plan must start with the announcement $\text{Malfunction}(m1, box)$ to the effect that the admins know of a malfunction, required for them to execute reboots. From s_0 , this results in s_1 of Figure 16. From there, both $\text{Reboot}(a2, sn1)$ and $\text{Reboot}(a2, sn2)$ must be performed before α_2 knows that no machine is malfunctioning. Hence, the

two shortest successful plans in which only a_2 reboots are

$$\begin{aligned}\pi' &:= \text{Malfunction}(m1, box) \otimes \text{Reboot}(a2, sn1) \otimes \text{Reboot}(a2, sn2) \\ \pi'' &:= \text{Malfunction}(m1, box) \otimes \text{Reboot}(a2, sn2) \otimes \text{Reboot}(a2, sn1)\end{aligned}$$

Hence, the potential uncertainty introduced by the non-rigidity of constants may be consequential for epistemic planning. In this simple example, the non-rigid serial numbers refers to objects, but as also agent terms may be non-rigid, the presented framework allows modeling situations in which e.g. an agent is uncertain about who a message must be delivered to, or planning in situations involving ‘code names’, known only to a strict subset of agents.

7 Languages for Actions

We define a language for reasoning about actions, denoted \mathcal{L}_{AM} . This language extends the basic language \mathcal{L} with *action modalities* with the form $[A, e]$, where A is an action model and e is an event from A . The language \mathcal{L}_{AM} has formulas of the form $[A, e]\varphi$, which are interpreted as: ‘after event e of action A occurs, φ is true’. This language extension allows us to include formulas mentioning other actions in the pre- and postconditions of some actions, as well as in goal formulas. It is thus possible to define, e.g., a goal such as: “Achieve a state in which it is impossible to perform an action that will result in φ ”. With finitely many actions described by the models $\mathcal{A} = \{A_1, \dots, A_n\}$, such a formula would be $\bigwedge_{A \in \mathcal{A}, e \in EA} [A, e] \neg \varphi$.

The grammar of \mathcal{L}_{AM} is defined by double recursion, adapting a construction well known in the DEL literature (see, e.g., appendix H in [8] or [28]).

Definition 15. Let $\mathcal{L}_0 = \mathcal{L}$, and let AM_0 be the set of pointed action models whose precondition formulas are all from \mathcal{L}_0 . Define \mathcal{L}_{k+1} and AM_{k+1} as follows:

$$\varphi ::= r(t_1, \dots, t_\ell) \mid \neg \varphi \mid \varphi \wedge \varphi \mid K_t \varphi \mid \forall x \varphi \mid [A, e] \varphi \quad (\mathcal{L}_{k+1})$$

where $(A, e) \in AM_k$, and let AM_{k+1} be the set of pointed action models whose precondition formulas are all from \mathcal{L}_{k+1} . Lastly, define the *language* \mathcal{L}_{AM} and the set of *action models* AM as

$$\mathcal{L}_{AM} := \bigcup_{k \in \mathbb{N}} \mathcal{L}_k, \quad AM := \bigcup_{k \in \mathbb{N}} (AM_k)$$

As with the formulas from the static language \mathcal{L} , the formulas from \mathcal{L}_{AM} are evaluated over epistemic models.

Definition 16. The satisfaction relation between epistemic models, assignments and formulas of \mathcal{L}_{AM} is the smallest extension of \models that satisfies:

$$M, w \models_v [A, e] \varphi \text{ iff } M, w \models_v \text{pre}(e) \text{ implies } M \otimes A, (w, e) \models_v \varphi$$

This extended satisfaction relation makes it possible to model-check conditions concerning actions. Given a pointed model (M, w) , we may want to know whether a formula φ would hold after a sequence of pointed action models $(A_1, e_1), \dots, (A_n, e_n)$ has been executed. This can of

First-order principles		Modal and interaction principles	
all propositional tautologies		$K_t(\varphi \rightarrow \psi) \rightarrow (K_t\varphi \rightarrow K_t\psi)$	K
$\forall x\varphi \rightarrow \varphi(y/x)$, for y free in φ	UE	$\forall xK_t\varphi \rightarrow K_t\forall x\varphi$, for x not occurring in t	BF
$t = t$, for $t \in \mathbb{T}$	Id	$(x \neq y) \rightarrow K_t(x \neq y)$	KNI
$(x = y) \rightarrow (\varphi(x) \rightarrow \varphi(y))$	PS		
$(c = c) \rightarrow \exists x(x = c)$	\exists Id		
$x \neq y$, if $\tau(x) \neq \tau(y)$	DD		
		Inference rules	
		From $\varphi, \varphi \rightarrow \psi$, infer ψ	MP
		From φ , infer $K_t\varphi$	KG
		From $\varphi \rightarrow \psi$, infer $\varphi \rightarrow \forall x\psi$, for x not free in φ	UG

Table 1: Axiom schemata for the minimal normal term-modal logic K.

course be done by computing a sequence of product updates and checking whether $M \otimes A_1 \otimes \dots \otimes A_n, (w, e_1, \dots, e_n) \models \varphi$. But, equivalently, we can check whether the corresponding formula holds at (M, w) , i.e., whether $M, w \models [A_1, e_1] \dots [A_n, e_n]\varphi$. If φ is a goal formula and $(A_1, e_1), \dots, (A_n, e_n)$ is a plan, then model-checking such a formula corresponds to *plan verification*. Section 8.2.2 gives so-called reduction axioms for \mathcal{L}_{AM} formulas, showing that any formula containing an action modality can be expressed as a formula in the basic epistemic language \mathcal{L} . Consequently, plan verification could be treated as a problem of model-checking formulas of \mathcal{L} in an initial state $s_0 = (M, w)$.

8 Axiomatic Systems and Metatheory

This section presents axiom systems for both static and dynamic term-modal logic. Metatheoretical results include soundness and completeness, frame characterizations, and decidability results. All proofs may be found in Appendix A.

8.1 Normal Term-Modal Logic

8.1.1 Axiom System

Table 1 contains the axioms and inference rules for the term-modal logic K. Some are common first-order axioms, like *Universal Elimination* (UE), *Reflexivity of Identity* (Id), and the *Principle of Substitution* (PS). In a modal logical context, PS also has a modal feature: it is restricted to variables to allow for non-rigid constants. If PS is assumed also for constants, $(a = b) \rightarrow (K_t\varphi(a) \rightarrow K_t\varphi(b))$ becomes a theorem, valid only for rigid constants. *Existence of Identicals* (\exists Id) is included to ensure that all constants obtains an extension in the canonical models of Section A.1.2; *Divided Domain* (DD) is included to enforce type-distinction between variables logically rather than syntactically. The modal and interaction principles *Distribution* (K) and *Knowledge of Non-Identity* (KNI) are formulated as standard while the *Barcan Formula* (BF) has a restriction in the term-modal case; the Barcan Formula ensures constant domains: its validity implies non-growing domains, illustrated in the proof of soundness (Section A.1.1), and its converse implies non-shrinking domains (and is provable in K, cf. e.g. [44, p. 245]). Knowledge of Non-Identity reflects the rigidity of variables. The inference rules *Modus Ponens* (MP), *Knowledge Generalization* (KG) and *Universal Generalization* (UG) contain no surprises.

Notice that nothing in the language or axioms of K specify the number of agents in the system. The number of agents emerges as a definable frame characteristic, cf. Section 8.1.4.

8.1.2 Normality

In Section 8.1.3, we formally state that K is complete with respect to the class of all frames. The axioms and inference rules sufficient for a complete system are close to standard axiomatizations of first-order modal logic, cf. e.g. [18, 32, 44]. We take the close-to-standard format of the K axioms to indicate the innocence of the term-modal extensions of the syntax and semantics. This is further corroborated by the main result of this section, the *Canonical Class Theorem* on this page. In essence, the theorem shows that any closed extension of K is complete with respect to the class of its canonical models. The result thus justifies the following definition:

Definition 17. A set of formulas $\Lambda \subseteq \mathcal{L}$ is called a *normal term-modal logic* if, and only if, Λ contains all axioms of Table 1 and is closed under the Table 1's inference rules. The smallest normal term-modal logic is denoted K .

8.1.3 Canonical Class Theorem and Completeness

In ordinary modal logic, each normal modal logic gives rise to a unique *canonical model*. In a similar manner, each normal term-modal logic Λ gives rise to a *class* of canonical models, one for each Λ -maximal consistent set. Section A contains the details of the construction, as well as the proof of the following main theorem:

Theorem 3 (Canonical Class Theorem). *Any normal term-modal logic Λ is strongly complete with respect to its canonical class.*

Mirroring the role of the Canonical Model Theorem of ordinary modal logic (see e.g. [12]), we obtain the following corollary to Theorem 3:

Corollary 1 (Completeness). *The logic K is strongly complete with respect to the class of all frames F .*

K is also sound with respect to the class of all frames. Section A.1.1 contains the formal statement and a proof sketch, with details given for the axiom K and the Barcan Formula.

8.1.4 Characterizing Frame Properties

The completeness result of Corollary 1 may be extended to more specific frame classes. Table 2 contains an overview of axiom schemata and the frame conditions they characterize. For illustration, proofs for 4 and N are given in Section A.1.3. From the Canonical Class Theorem and Table 2, completeness results for standard logics like $KD45$, $S4$ and $S5$ follow as corollaries.




Figure 17: A transitive model invalidating $K_c(b = c) \rightarrow K_c K_c(b = c)$ at w . With $D_{\text{agt}} = \{\alpha, \beta\}$, all relations are transitive. The notation $a \mapsto \alpha$ specifies that α is the extension of the constant a in the given world. In w , it holds that $K_c(b = c)$ as $c \mapsto \alpha$ in w and $b, c \mapsto \beta$ in w' (in w , the knowledge that $(b = c)$ is held by agent α , as $c \mapsto \alpha$ in w). World w' satisfies $K_c(b \neq c)$ as $c \mapsto \beta$ in w' and $b \mapsto \beta, c \mapsto \alpha$ in w'' (in w' , the knowledge that $(b \neq c)$ is held by β , as $c \mapsto \beta$ in w'). As a consequence, w' also satisfies $\neg K_c(b = c)$. Hence, w does not satisfy $K_c K_c(b = c)$ —that the agent named c (i.e., α) knows that the agent named c (i.e., β) knows that $(b = c)$.

Axiom		Frame condition
$\forall x (K_x \varphi \rightarrow \varphi)$	T	Reflexive
$\forall x (\neg K_x \perp)$	D	Serial
$\forall x (K_x \varphi \rightarrow K_x K_x \varphi)$	4	Transitive
$\forall x (\neg K_x \varphi \rightarrow K_x \neg K_x \varphi)$	5	Euclidean
$\exists x_1, \dots, x_n \left(\left(\bigwedge_{i \leq n} K_{x_i} \top \right) \wedge \left(\bigwedge_{i, j \leq n, i \neq j} x_i \neq x_j \right) \wedge \forall y (K_y \top \rightarrow \bigvee_{i \leq n} y = x_i) \right)$	N	$ D_{\text{agt}} = n$
$\exists x_1, \dots, x_m \left(\left(\bigwedge_{i, j \leq m, i \neq j} x_i \neq x_j \right) \wedge \forall y \left(\bigvee_{i \leq m} y = x_i \right) \right)$	M	$ D = m$

Table 2: Term-modal axiom schemata and the frame conditions they characterize.

The principles N and M are special to our term-modal treatment. N and M define domain sizes. Nothing in the language or axioms of K specify the number of agents in the system: as in first-order logic, the domain size is by default left unspecified. In ordinary epistemic logic, it is common to assume a fixed, finite index set of agents. The domain size principles N and M similarly fixes domain sizes: N fixes the agent domain to size n . It uses the $K_x \top$ -expressions to ensure that the bound variables are all of type agt. With the quantifications thus ranging only over agents, N specifies specifically the size of D_{agt} . This may be compared to M, which does not put constraints on the type of the bound variables, thereby fixing only the size of the joint agent–object domain D . For details concerning N, see Proposition 6 on 44.

The principles T, D, 4 and 5 deviate from their ordinary forms by being quantified. In standard modal logic, the formula

$$K_c \varphi \rightarrow K_c K_c \varphi \quad (1)$$

characterizes the class of transitive frames. This is not true here, as the constant c may be *non-rigid*:¹ see Figure 17 for a transitive model invalidating (1). The invalidity arises as the extension of the c is not fixed under scope of operators: in the consequent, the accessibility relation which the inner occurrence of K_c quantifies over need not be the same as the accessibility relation of the outer. This makes the appeal to transitivity void.² The formulation in Table 2 avoids the non-rigidity problem,

¹The non-rigidity of constants is reflected in K: *Knowledge of Identity* is provable for variables, but not for constants. I.e., K proves $(x = y) \rightarrow K_t(x = y)$, but not $(a = b) \rightarrow K_t(a = b)$.

²In his 1962 [40], Hintikka argues that $K_c \varphi \rightarrow K_c K_c \varphi$ intuitively is valid only if c knows that she is c ; i.e., that she knows who c is. Hintikka argues that this is captured by $\exists x K_c(x = c)$, which makes c locally rigid for the agent: $I(c, w') = I(c, w)$ for

but does impose the criteria for all agents uniformly.

8.1.5 Heterogeneous Agents

Though treating all agents uniformly is common in epistemic logic, one may desire heterogeneous agents. With the given setup, we do not believe this can be done at the level of frames. On the level of models, one option to this end is to attribute epistemic criteria to subgroups using predicates; a second is to introduce individual names. In either case, one may desire the defining criterion to be rigid. However, full rigidity is not definable in general as models may be disconnected. *Local rigidity*—invariance of interpretation over connected components—is definable by formulas of the forms

$$\forall x(r(x) \leftrightarrow \forall yK_y r(x)) \quad (2)$$

$$\exists x((x = a) \wedge \forall yK_y(x = a)) \quad (3)$$

The validity of (2) and (3) characterize features of interpretations: (2) ((3), resp.) is valid in a model $M = (D, W, R, I)$ iff for all $w, w' \in W$, $(w, w') \in R_\alpha$ for some $\alpha \in D_{\text{agt}}$ implies $I(r, w) = I(r, w')$ ($I(a, w) = I(a, w')$, resp.). In conjunction with formulas of the forms

$$\forall x(r(x) \rightarrow (K_x \varphi \rightarrow K_x K_x \varphi)) \quad (4)$$

$$K_a \varphi \rightarrow K_a K_a \varphi \quad (5)$$

one may obtain some individuated control over relation properties.

8.1.6 Decidability

Let K_n and $K_{n/m}$ be the smallest normal extensions of K with, respectively, the domain size axiom N , and both domain size axioms N and M under the proviso that $m > n$. K_n and $K_{n/m}$ are then sound and complete with respect to, respectively, the class of all frames with exactly n agents, and the class of all frames with exactly n agents and exactly $m - n$ objects. These finite domain properties are used in the proof of items 1. and 2. of the below proposition, shown in Section A.1.4. Decidability results from the literature are discussed in Section 9.3.

Proposition 1. *Let $K_{n/m}$, K_n and K be given in \mathcal{L} , based on the signature Σ . Let $\mathcal{L}_{\text{agt}} \subseteq \mathcal{L}$ contain all formulas containing only agent-terms, $t \in \mathfrak{t}^{-1}(\text{agt})$.*

1. For all $\varphi \in \mathcal{L}$, it is decidable whether $\vdash_{K_{n/m}} \varphi$ or not.
2. a) For all $\varphi \in \mathcal{L}_{\text{agt}}$, it is decidable whether $\vdash_{K_n} \varphi$ or not. b) In general, it is undecidable.
3. In general, it is undecidable whether $\vdash_K \varphi$ or not.

8.2 Dynamic Term-Modal Logic

8.2.1 Axiom System

Table 3 contains the axioms and inference rules for the dynamic term-modal logic AM. In Section 8.2.2, we formally state that $K + \text{AM}$ is sound and complete with respect to the class of all frames.

all w' in $R_{I(c,w)}(w)$. Indeed, $\exists xK_c(x = c) \wedge K_c \varphi \rightarrow K_c K_c \varphi$ is valid on transitive frames.

This completeness result may be extended to more specific frame classes, as was the case with K (see Section 8.1.4). The completeness proof for $K + AM$ is by translation, a well-known approach in DEL [7, 8, 28, 72]. The axioms in AM are so-called *reduction axioms*, which enable the translation of formulas with action modalities into provably equivalent ones without any action modalities. Then completeness follows from the known completeness of the static logic K . For a detailed description of the reduction strategy to completeness, see e.g. [28].

Reduction axioms

$[A, e]r(t_1, \dots, t_n) \leftrightarrow (\text{pre}(e) \rightarrow \text{post}(e)(r(t_1, \dots, t_n)))$	Action and atom
$[A, e]\neg\varphi \leftrightarrow (\text{pre}(e) \rightarrow \neg[A, e]\varphi)$	Action and negation
$[A, e](\varphi \wedge \psi) \leftrightarrow (([A, e]\varphi) \wedge ([A, e]\psi))$	Action and conjunction
$[A, e]K_t\varphi \leftrightarrow \bigwedge_{e' \in E} (Q(e, e')[x^* \mapsto t] \rightarrow K_t[A, e']\varphi)$	Action and knowledge
$[A, e]\forall x\varphi \leftrightarrow (\text{pre}(e) \rightarrow \forall x[A, e]\varphi)$	Action and quantification (Dynamic Barcan)
$[A, e][A', e']\varphi \leftrightarrow [(A, e) \circ (A', e')]\varphi$	Action composition

Inference rules

From φ , infer $[A, e]\varphi$	Action necessitation
--	----------------------

Table 3: Axiom and rule schemata for the system AM .

The reduction axioms in AM are similar to those used in logics for epistemic actions, introduced by [7]. Naturally, as dynamic term-modal logic is first-order, there are reduction axioms for formulas involving quantifiers. Moreover, the axiom for formulas with the knowledge operator is non-standard. Unlike standard action models, the ones presented here are edge-conditioned and use variable substitutions, which require some modifications. A more detailed comparison of these axioms and standard ones is provided in Section 9.2. The *Action composition* axiom appeals to action models of the form $(A, e) \circ (A', e')$. This notation refers to the *composition* of (A, e) and (A', e') , defined following [30], but adapted to accommodate edge-conditions and first-order atoms:

Definition 18. Let $A_1 = (E_1, Q_1, \text{pre}_1, \text{post}_1)$ and $A = (E_2, Q_2, \text{pre}_2, \text{post}_2)$ be given. The *composition* of A_1 and A_2 is the action model $A_1 \circ A_2 = (E, Q, \text{pre}, \text{post})$ where

1. $E = E_1 \times E_2$
2. $Q((e_1, f_1), (e_2, f_2)) = Q_1(e_1, f_1) \wedge [A_1, e_1]Q_2(e_2, f_2)$.
3. $\text{pre}(e_1, e_2) = \text{pre}(e_1) \wedge [A_1, e_1]\text{pre}(e_2)$,
4. $\text{dom}(\text{post}((e_1, e_2))) = \text{dom}(\text{post}_1(e_1)) \cup \text{dom}(\text{post}_2(e_2))$ and if $r(t_1, \dots, t_k) \in \text{dom}(\text{post}((e_1, e_2)))$, then

$$\text{post}((e_1, e_2)(r(t_1, \dots, t_k))) = \begin{cases} \text{post}_1(e_1)(r(t_1, \dots, t_k)) & \text{if } r(t_1, \dots, t_k) \notin \text{dom}(\text{post}_2(e_2)) \\ [A, e]\text{post}_2(e_2)(r(t_1, \dots, t_k)) & \text{otherwise} \end{cases}$$

8.2.2 Soundness and Completeness via Reduction Axioms

As anticipated in Section 8.2.1, $K + AM$ is sound and complete with respect to the class of all models.

Proposition 2 (Soundness of $K + AM$). $K + AM$ is sound with respect to the class of all models.

The soundness of $K + AM$ (Proposition 2) is established in the standard way, by showing the semantic validity of the reduction axioms and inference rules. The proof is straightforward and therefore omitted.

Completeness follows as a corollary from a number of lemmas that are presented in Section A.2.

Corollary 2 (Completeness of $K + AM$). $K + AM$ is complete with respect to the class of all models. Moreover, any extension of $K + AM$ obtained by adding axioms characterizing frame conditions is complete with respect to the corresponding class of models.

9 Related Work

9.1 Epistemic Planning

Several articles on multi-agent epistemic planning have appeared recently. The existing work can be organised along the following categories: modeling of epistemic planning, tractability and complexity, and implementation and applications.

On the modeling side, multiple articles have presented formalisms for multi-agent epistemic planning based on DEL [2, 15, 58, 92]. These models are very expressive, capturing several key aspects of multi-agent epistemic planning. These aspects include: epistemic actions and goals, higher-order knowledge and belief, partial observability, etc. A thorough comparison of the present framework with existing DEL formalisms is found in Section 9.2.

The rich expressivity of DEL comes at a cost, as planning problems specified in DEL are in general computationally difficult to solve (more on this below). This has partly motivated the introduction of simpler formalisms for epistemic planning. Some of these formalisms build on classical planning. The model in [71] extends STRIPS to allow knowledge declarations in preconditions and postconditions. The framework is however restricted to single-agent planning, does not support higher-order reasoning, and allows only a restricted form of quantification. The multi-agent planning frameworks in [46, 61] follow a compilation approach, translating restricted fragments of epistemic planning into classical planning languages.

The approaches in [42, 45] describe planning domains via a type of state-transition system extended with epistemic information, called a *concurrent epistemic game structure* (CEGS). This representation makes it easy to define multi-agent notions such as ‘joint action’ or ‘multi-agent plan’. However, the representation inherits some of the well-known problems of transition-system models, including the lack of compact descriptions of actions and efficient heuristics that can avoid building the full state-transition system when planning (see [14] for a discussion of these and other limitations).

The non-DEL formalism that most closely resembles the approach of this paper is the epistemic game description language GDL-III [86]. The language is epistemic and first-order. A key feature of this language is that only what agents can see and do has to be defined. This is done via declarations that use the keywords **Sees** and **Does**, which loosely correspond to modalities. GDL-III has a simple syntax and allows compact specifications of actions. For instance, the following GDL-III rules [9] describe schematically communication actions which are similar to the ones from Example 5:

$$\begin{aligned} \mathbf{Sees}(x, i_a) &\Leftrightarrow \mathbf{Does}(i, \text{announce}(z)) \wedge \text{Obs}(i, x) \\ \mathbf{Sees}(y, \varphi) &\Leftrightarrow \mathbf{Does}(i, \text{announce}(\varphi)) \wedge \text{Obs}(i, y) \wedge \text{Listen}(i, y) \wedge \varphi \end{aligned}$$

These two transition rules are interpreted as follows: if agent i announces z then any agent x observing i will receive the information i_a . Any agent y that observes i and listens to i will learn the content of the announcement, φ . Given the semantics of GDL-III, it follows that agents who only see i_a will know that i made an announcement but will not learn the content of the announcement. Agents who observe i , however, will know that φ must be true. Moreover, if an agent x observes agent i , does not listen to i , but knows that another agent y listens to i , then the semantics entails that x will know that y will know the content of the announcement after it has been made. The model is therefore schematic and context-sensitive, like the epistemic action schemas presented here. The syntax of GDL-III is simpler than that of DEL when it comes to representing actions. However, as noted in [31], specifying nested and higher-order knowledge is more difficult in GDL-III than in DEL, and the formalism requires more involved semantics. The work in [31] provides a detailed comparison of DEL and GDL-III, concluding that GDL-III offers a simpler syntax, while DEL provides simpler semantics. In [31], it is shown that large fragments of GDL-III and DEL are equally expressive by giving compilations between the two.

Concerning decidability and complexity, it was first shown in [15] that the general plan existence problem in propositional DEL planning (i.e., deciding whether a plan exists given a multi-agent planning task) is undecidable. In fact, the problem is undecidable with two agents only, no common knowledge, and no postconditions. In [52] it is shown that public actions are enough for undecidability when the initial state meets certain technical conditions. That paper also identifies an undecidable subclass of small epistemic planning problems comprising two agents, one action, six propositions and a fixed goal. The undecidability results straightforwardly apply also to the present framework.

Although the general problem is undecidable, several papers have identified decidable fragments of epistemic planning that are still reasonably expressive. Single-agent epistemic planning is decidable [15]. The multi-agent problem becomes decidable if actions are only allowed to have propositional preconditions (i.e., no epistemic formulas appear in the preconditions) [92]. The computational complexity of this fragment belongs to $(d + 1)$ -EXPTIME for a goal formula whose modal depth is d . If actions are restricted to have propositional preconditions and no postconditions, the plan existence problem becomes PSPACE-complete [21]. Stronger restrictions, such as allowing only private and public announcements, bring the complexity down to NP-complete [17]. As mentioned in Section 6, in [55] we show that single-agent epistemic planning and multi-agent planning with non-modal preconditions are also decidable in the term-modal, first-order case described in this paper, echoing the results for propositional planning in [15, 92].

As for implementation and applications, a number of techniques and planners have been developed over the last decade. An approach that has gained popularity is the compilation approach. The idea involves choosing a suitably restricted fragment of DEL that can be encoded in a classical planning language. Epistemic problems are then translated into classical ones so that state-of-the-art planners can be used to solve them efficiently. The compilations rely on different restrictions. The system in [46] assumes that actions are public, physical actions are deterministic, and that all agents start with a common initial belief on the set of worlds that are possible. The paper adopts a centralised perspective, with planning done off-line from the viewpoint of a single agent. In [47], the authors extend this framework to cover on-line planning from the perspective of the agents themselves. The planner in Muise *et al.* [61] requires a finite depth of nesting of modalities and no disjunctions. Cooper *et al.* [23] use an encoding based on special variables describing what agents can see. The epistemic problems expressible with this restricted language are then encoded in PDDL and solved using the Fast Downward planner [39]. As mentioned before, the PKS system in [71] encodes epistemic planning using a STRIPS-like language. This language can describe single-agent, epistemic planning problems

with conditional effects. The PKS system tries to solve these problems using an efficient but incomplete algorithm.

A small number of epistemic planners do not rely on compilation into classical planning. The system MEPK [43] performs multi-agent epistemic planning from the third-person viewpoint. The system can handle private actions and beliefs, as formalized with the modal logic KD45n. The system does not support arbitrary common knowledge but can deal with a weaker form of common knowledge. Finally, Le *et al.* [51] present two forward planners, called EFP and PG-EFP, for multi-agent epistemic planning. These planners can deal with unlimited nested beliefs, common knowledge, and epistemic goals when the number of worlds in the initial state is not too large.

9.2 Dynamic Epistemic Logic

There is a vast and excellent literature on both epistemic logic and dynamic epistemic logic to which the reader is referred for both technical and conceptual introductions—see e.g. [8, 10, 28, 32, 40, 60, 53].

The approach to modeling actions taken in this paper is based on the idea of *action models* applied using *product update* as introduced first by Baltag, Moss and Solecki [7]. The reduction axiom approach to proving completeness for logics with actions was first suggested by Plaza [72] for the case of *truthful public announcements*. Our approach is the same, but for general actions models. It is based on [6, 8, 28].

The format of the action models presented here differs mainly in four aspects from those introduced in [7]: our action models have *postconditions*; are *first-order* rather than propositional; accommodate *term-modal relations*; and have *conditioned edges*.

Our approach to postconditions is inspired by [11, 30]. From there, it is a straightforward generalization to alter pre- and postconditions to allow updates of first-order Kripke models.

A substantial departure from the standard is the accommodation of term-modal relations and the edge-conditioning. The definition avoids two problems for term-modal action models—one pointed out by Kooi [48] and one concerning reduction axioms—by an adjustment of the propositional *edge-conditioned* action models of Bolander [13].³

In the standard definition, an action model A for index set of agents \mathcal{I} consists of a finite set of events $E = \{e, \dots, e'\}$ and a map $R : \mathcal{I} \rightarrow \mathcal{P}(E \times E)$, plus assignments of pre- and postconditions. In the term-modal treatment, the set \mathcal{I} is a proper part of the semantics of state representations. Adding an operator $[A, e]$ to the language thus conflates syntax and semantics, Kooi points out.

In considering reduction actions, we found that this problem runs deep. Consider the standard reduction axiom for the modal operator:

$$[A, e]K_i\varphi \leftrightarrow \left(\text{pre}(e) \rightarrow \bigwedge_{f:(e,f) \in R_i} K_i[A, f]\varphi \right) \quad (6)$$

In (6), the agent index i links the occurrences of the modal operator K_i with the relation R_i used in the quantifying conjunction. This link is broken in the term-modal treatment: the “ i ” indexing the operators is a syntactic term, while the “ i ” indexing the relation is an element of a domain of quantification. Without consulting an interpretation (or variable assignment), these two occurrences are unlinked: there is no guarantee that R_i is the relation used in evaluating $K_i\varphi$.

³Both approaches result in context-sensitive actions: the distinguishability of two events depends on model to be updated. See [13, 53, 77, 78, 79] for arguments to the effect that more context-sensitivity than what is present in standard action models is desirable.

To resolve the conflation problem, Kooi defines action models where accessibility relations over events E are assigned to *groups* of agents on a per-application basis: With Φ a finite set of mutually inconsistent and jointly exhaustive formulas with free variable x , each pointed model (M, w) and variable valuation v defines a partition on the agent domain with cells $\{d \in D_{\text{agt}} : M, w \models_{v[x/d]} \varphi(x)\}$ for each $\varphi(x) \in \Phi$; each agent in such a cell (group) is assigned the same accessibility relation using a map $S : \Phi \rightarrow \mathcal{P}(E \times E)$. In effect, the action model makes no direct reference to the agent domain, thus avoiding the conflation problem.

Additionally, Kooi's definition yields a solution to the problem of unlinked indices as the relations of the action model may now be referred to using syntactical constructs. With $\Phi = \{P_1(x), \dots, P_n(x)\}$, a suggestion for a reduction axiom could be

$$[A, e]K_t \varphi \leftrightarrow \left(\text{pre}(e) \rightarrow \bigwedge_{k \leq n} \left(P_k(t) \rightarrow \bigwedge_{f: (e, f) \in R(P_k)} K_t[A, f] \varphi \right) \right).$$

We obtain a similar solution by adjusting the edge-conditioned action model of Bolander [13]. In an edge-conditioned action model, whether two events are related for some agent $i \in \mathcal{I}$ is conditional on whether a given formula is satisfied in the pointed model on which the action is executed. Formally, each agent-edge pair is assigned a condition by a map $Q : \mathcal{I} \rightarrow (E \times E \rightarrow \mathcal{L})$.

Inspired by both Bolander and Kooi, we use a map $Q : E \times E \rightarrow \mathcal{L}$ where $Q(e, e')(x^*)$ has exactly one free variable, x^* . When the resulting action model is executed on a pointed model (M, w) , an edge is present for an agent $\alpha \in D_{\text{agt}}$ if $M, w \models_{v[x^* \mapsto \alpha]} Q(e, e')(x^*)$. As the condition $Q(e, e')(x^*)$ is a formula, this approach allows the formulation of reduction axioms, cf. Section 8.2.1.

Our version of the Q function and Kooi's approach S are equally general. Given an action model $(E, S, \text{pre}, \text{post})$ with $S : \Phi \rightarrow \mathcal{P}(E \times E)$, let $Q : E \times E \rightarrow \mathcal{L}$ be given by $Q(e, e') = \varphi$ such that $(e, e') \in S(\varphi)$. Then Q emulates S : for all models M , $M \otimes (E, S, \text{pre}, \text{post}) = M \otimes (E, Q, \text{pre}, \text{post})$. *Vice versa*, to emulate a map Q , for each $A \subseteq E \times E$, let

$$\varphi_A := \bigwedge_{\varphi \in Q(A)} \varphi \wedge \bigwedge_{\psi \in \Psi} \psi \tag{7}$$

with Ψ the largest subset of $\{\neg\varphi : \varphi \in Q(E \times E) \setminus Q(A)\}$ such that (7) is consistent. Then $S : \varphi_A \mapsto A$ for each $A \subseteq E \times E$ is a Kooi map that emulates Q . We opt for the edge-conditioned formulation due to its correspondence with the standard precondition maps $\text{pre} : E \rightarrow \mathcal{L}$.

Finally, note that both may emulate standard action models over classes of models where each agent α is designated by a rigid constant a_α (as is conceptually implied by identifying agents with indices). The standard map $R : D_{\text{agt}} \rightarrow \mathcal{P}(E \times E)$ may be emulated by the map $Q : E \times E \rightarrow \mathcal{L}$ with $Q(e, e') = \bigvee_{a_\alpha : (e, e') \in R_\alpha} (x = a_\alpha)$.

9.3 Term-Modal Logic

The term-modal treatment of epistemic operators as behaving both as modal operators and as first-order predicates was suggested already by von Wright in his 1951 [88], though the direction was not formally explored. Formally, Hintikka allowed the constructions in his 1962 [40], and the term-modal aspects are used in discussions concerning the validity of $K_a \varphi \rightarrow K_a K_a \varphi$, where Hintikka notes that the schema is only valid if *a knows who a is*, captured by $\exists x K_a (x = a)$ (see also Section

8.1.4 on frame characterizations). Semantically, Hintikka linked individuals and operators in [41] using world-relative first-order interpretations extended to assign alternatives to individuals in the domain of quantification, D . Work in philosophical logic followed Hintikka’s term-modal syntax—even called “standard” by Carlson in 1988 [20]—but the semantic link did not pertain: [84] exemplifies a pseudo-use. Carlson enforced the semantic link, using a partial map $R : D \rightarrow \mathcal{P}(W \times W)$ to assign accessibility relations to individuals. He further presents a Hintikka-style model set proof theory for a three-valued Kripke-style semantics with non-rigid terms, varying domains and reflexive relations, and shows completeness.

In computer science, a format similar to Carlson’s is frequently used when giving the semantics for propositional epistemic logic, with the set of agents D treated as an *index set* instead of a domain of quantification, even in the first-order case: e.g., in Fagin *et al.*’s first-order treatment [32], in a formula like $K_{Alice} Governor(California, Pete)$, both *California* and *Pete* are first-order terms, but *Alice* is not—*Alice* is an agent. Here, then, agents and their names are *equated*.

The issue of equating agents and their names, and why this is unsatisfactory in many computer science applications, is discussed at length by Grove & Halpern [38] and Grove [37]. They identify the following inadequacies: systems that equate agents with their names cannot represent agent sets of non-fixed size, do not allow for reference to agent groups, for non-rigid names, nor for indexical and relative reasoning (using terms like “me” to express e.g. “the agent to the left of me”). In response, [38] develops a propositional epistemic logic with indexical reference obtained by evaluating formulas at agent-world pairs, on which [37] builds a first-order variant to additionally handle issues of *de dicto/de re*-like reference scope, as well as multi-naming of agents. The latter is in effect a variant of non-rigid constant, varying domain, term-modal logic with formulas evaluated at agent-world pairs. Further, the language contains two sorts, one for *agents* (like our agt terms), and one for *names*.⁴ This allows explicit reasoning about *naming*. Adding a third sort to the present framework would be unproblematic, but the indexical semantics would require in-depth re-working. Similarly would varying agent domains require work, unless emulated by an existence predicate, cf. [33]. Beyond this, the present framework tackles the issues raised in [37, 38]: agents and names are not equated by the use of (non-rigid) constants of sort agt, that additionally allow for multi-naming; agents groups may be denoted by predicates and relative properties by relations; finally, *de dicto/de re* distinctions are expressible using quantification. However, beyond the formal difference and similarities, we would find an in-depth philosophical comparison of the interpretation of the two frameworks interesting. In [54], we illustrate the system presented here with examples that touch on several of the involved issues.

One reason for sticking with ordinary modal operators even in a first-order setting is that term-modal operators adds design choices and possible complications, as discussed by Lomuscio & Colom-betti in their early contribution to the term-modal literature [57]. In constructing a term-modal extension of multi-agent KD45 with non-rigid terms, they discuss how to evaluate formulas $B_a\varphi$ when a is not an agent denoting term. Intuitively, $B_a\varphi$ should be false, as only agents can truly hold beliefs, but—they remark—this would imply the invalidity of $B_a(\varphi \vee \neg\varphi)$. They conclude against a two-sorted approach, as a similar problem surfaces for formulas $B_aB_b\varphi$ when agent a believes that the term b denotes a non-agent.⁵ Ultimately, Lomuscio & Colom-betti opt for a *partial logic* with *truth-value gaps*, letting the truth-value of $B_a\varphi$ be *undefined* when a denotes a non-agent; they take a *valid* formula to be *sometimes satisfied*, but *never false*. Their semantics are constant domain, and each element is, at

⁴Such two sorts are also used by Rendsvig in a quantified, but not term-modal, epistemic logic analysis of semantic competence in relation to Frege’s puzzle about identity [75, 76].

⁵This obstacle is avoided in the present paper by syntactically forcing all operator-subscripts to be of the agent-sort.

each world, assigned a set of *doxastic alternatives*; an element is an *agent in world* w if it is assigned a non-empty set. Hence, agenthood is *world-relative*. They present an axiom system—which includes a term-modal Barcan formula $\forall y(B_x \varphi(y)) \rightarrow B_x \forall y(\varphi(y))$ and quantified frame-characterizing formulas like $\forall x(B_x \varphi \rightarrow B_x B_x \varphi)$ like the present paper—and show soundness, citing [56] for details.

Bivalent systems are presented by Thalmann [85] and Fitting, Thalmann & Voronkov [34], with these two works coining the label ‘term-modal logic’. In their setting, each world w is associated with an *inner domain* $D(w)$ of objects existing at w , with $D(w)$ a subset of the *outer domain* D , for all w . The inner domains are assumed *increasing*: if $wR_d w'$ for some $d \in D$, then $D(w) \subseteq D(w')$. Further, terms are assumed rigid and with an interpretation defined at every world ($I(c) \in D(w)$ for all $w \in W$). This combination seemingly⁶ eliminates the need for truth-value gaps, but the problems raised by non-agents are not discussed. For several classic frame-conditions, [34, 85] presents both sequent and tableau proof systems (K, D, T, K4, D4, S4).

Orlandelli & Corsi [63] also investigate sequent calculi for term-modal logics. Their semantics is more general as they omit the increasing domain requirement, and—as they also consider Euclidean frames—they also obtain completeness for more frame classes. The syntax is without constants, so the rigidity/non-rigidity dichotomy is non-applicable. The semantics are bivalent. The combination of varying domains and bivalent semantics is facilitated by the atomic formula satisfaction clause

$$M, w \models_v r(x_1, \dots, x_n) \text{ iff } (v(x_1), \dots, v(x_n)) \in I(r, w),$$

with $I(r, w) \subseteq D^n$ again with D the outer domain. E.g., with $I(=, w) = \{(d, d) \in D^2 : d \in D\}$, the formula $(x = x)$ is satisfied in (M, w) even if $v(x) \notin w$. However, as the quantifiers only range over the inner domain of worlds, the semantics oddly make $p(x) \wedge \forall y \neg p(y)$ satisfiable.

In [48], Kooi introduces a *dynamic* term-modal logic, including a first-ever first-order version of DEL action models. The language of [48] is first-order dynamic logic with wildcard assignment, but where the set of first-order terms is also the set of atomic programs, the models for which are constant agents-only domain with non-rigid terms (and very similar to our general case, but restricted to agents-only). This language is more expressive than ordinary term-modal logic. The first-order dynamic logic aspect implies that the validity problem is Π_1^1 complete, eliminating hope for a finitary proof system. However, the expressivity of the language allows the definition of a *non-rigid common knowledge*. If not for our two-sorted domain, our language and semantics could be seen as a special case of Kooi’s. Kooi’s action models are discussed in the next section.

Seligman & Wang [89] investigate a fragment Kooi’s system. The fragment allows only basic assignment modalities to form a quantifier-free term-modal logic (without function symbols), a fragment rich enough to express *de dicto/de re* distinctions and *knowing who* constructions in a setting where names are not common knowledge. The main result is a complete axiomatization for the class of S5 models. As Barcan-like formulas are not included in the investigated language fragment but are the common characterizers of constant domain semantics, this result is quite non-standard. The authors also discuss decidability: providing no hard results, they conclude “We are not that far from the decidability boundary, if not on the wrong side.”

Corsi & Orlandelli [24] introduce a generalization of term-modal syntax to be able to express the difference between *de dicto* and *de re* statements without invoking quantifiers. They introduce complex term-modal operators $|t: \overset{c}{x}|p(x)$ with the reading that t knows of c that (s)he is $p(x)$. These are

⁶Seemingly, as we are confused about the satisfaction clause for atomic formulas [34, Def. 7, It. 1], stating that $w, V \models R(t_1, \dots, t_n)$ iff $w \models R(V(t_1), \dots, V(t_n))$ with $V(t_i) \in D$, but no specification of the conditions for the right-hand condition, nor any specification of how the relation symbol R is assigned extension. However, if this is assumed settled as ordinarily (as in the present paper), the increasing domain assumption seems sufficient to obtain a well-behaved semantics, as is the case in ordinary first-order modal logic. See e.g. [35] for an introduction and [44] for details.

interpreted over so-called *epistemic transition structures with double-domains*. The resulting *indexed epistemic logics* are further investigated in [26, 25]. It would be interesting to know what the relationship is to the also expressive language of Kooi [48].

Where the domain of Kooi [48] consists only of agents, Rendsvig [74] introduces a model with a single-sorted language with non-rigid terms that denote elements in a constant domain containing both agents and objects. As in [57], this requires an *ad hoc* solution to the semantics of formulas $K_a\varphi$ when a denotes a non-agent. The solution used is to then interpret $K_a\varphi$ as a *global* modality. This preserves the bivalence of the systems while making all operators normal. As a result, [74] presents a canonical model theorem, facilitating completeness proofs for classic frame classes.

The semantics of this paper are based on Achen's [1], which in turn is a two-sorted refinement of [74]. What we consider an improvement of [1] over [74] is exactly the two-sorted approach: distinguishing between agent and object terms removes the need to define *ad hoc* semantics for knowledge operators indexed by non-agents. Taking a two-sorted approach eliminates the possibility of modeling agents that are uncertain about whether a given term refers to an agent or an object, but results in a system which we consider well-behaved.

Term-modal like, Naumov & Tao [62] present a propositional term-modal logic, but where operators may be indexed by sets of terms, making $\exists xK_{\{x,a\}}\varphi$ a formula. Such operators are given a *distributed knowledge* semantics in S5 models with constant agents-only domain and rigid terms for which a complete axiom systems is presented.

Sawasaki, Sano and Yamada [81] consider a term-modal syntax where operators are indexed by a sequence of terms making e.g. $\forall x\forall yK_{[x,y]}R(x,y)$ well-formed, with the intended deontic reading that x is obliged by y to ensure $R(x,y)$. They present complete axiom system and sequent calculi.

Sedlar [82] also uses a rigid terms, agents-only constant domain semantics to represent an epistemic logic of evidence using a term-modal language as that presented here. Sedlar shows that his term-modal framework is able to emulate monotonic modal logics and epistemic logics with awareness, obtaining a decidability result for the fragment with no constants nor functions, but 0-ary predicates and single unary predicate.

Several other authors have also looked at decidability issues for varieties of term-modal logics. Kooi [48] points out that the monadic fragment of his system is undecidable by a result of Kripke [50]. As Kripke's result concerns first-order modal logic in general (see e.g. [44, p. 271 ff.]), it applies to broadly to term-modal logics, too. For term-modal logics, Padmanabha & Ramanujam [67] even show that the propositional fragment is undecidable. As decidable, they identify the monodic fragment (formulas using only one free variable in the scope of a modality). [65] considers model checking for the fragment over a restricted model class and [64] presents a translation of the monodic fragment (without identity) into FOML.

In [66], Padmanabha & Ramanujam further investigate a variable-free propositional bi-modal logic with implicit quantification, with formulas $[\forall]\varphi$ and $[\exists]\varphi$ asserting that along all (resp. some) accessibility relations φ is necessary. These variable-free formulas thus correspond to the propositional term-modal formulas $\forall xK_x\varphi$ and $\exists xK_x\varphi$. The relevant logic is shown decidable, to be bisimulation-invariant fragment of an appropriate two-sorted first-order logic, related to the 'bundled fragment' of term-modal logic. Model checking for the system is investigated in [69]. In [68] Padmanabha & Ramanujam, turn to the two variable fragment of term-modal logic, which they show decidable. The thesis [64] collects a selection of the mentioned results, and additionally presents a translation of TML without identity into propositional TML.

For their own system, Orlandelli & Corsi [63] show two fragments decidable, the first propositional with quantifiers and operators occurring only in pairs of the forms $\exists x[x]$ or $\forall x(x)$. This

fragment simulates non-normal monotone epistemic logics. The second fragment allows expressing 1-ary groups' higher-order knowledge about proposition symbols, e.g. with $\forall x(p(x) \rightarrow K_x(K_y q))$ an allowed formula. Also Pliuškevičius & Pliuškevičienė [73] treats a fragment of propositional term-modal logic, but with, term-modal operators for belief and mutual belief, allowing only pairwise quantifier-operator nestings (e.g., for p a propositional atom, $\forall x B_x \exists y B_y p$ is well-formed, while $\forall x \exists y B_x B_y p$ is not). For their agents-only constant domain KD45 semantics, they present a terminating sequent calculus decision procedure. For further decidability results, it may be relevant to consult Shtakser [83], who investigates propositional modal languages includes quantification over modal operators and predicate symbols that take modal operators as arguments.

Beyond its main decidability result, Padmanabha & Ramanujam [68] also discusses translation of term-modal logic into first-order modal logic. In a setting with no constants or function symbols, the authors suggest a translation of TML into FOML with a single modality K and a new unary predicate P , inductively translating $K_x \varphi$ to $K(P(x) \rightarrow \varphi)$ and $\hat{K}_x \varphi$ to $\hat{K}(P(x) \wedge \varphi)$. [68] omits the details, but claims this translation produces FOML formulas equi-satisfiable with their TML originals. This suggests that completeness results for term-modal logics may also be shown indirectly via translation and application of well-known results for FOML (see e.g. [44]), instead of by the direct constructions found in the Appendix.⁷ Whether a translation approach would work for the present framework is an open question, but we have reservations concerning the general applicability of the suggested translation.⁸

10 Final Remarks

We conclude with open questions we see in relation to epistemic planning, and a summary of the main contributions of the paper. The following are some possible avenues for future research:

1. *Decidability and complexity.* As presented in the literature review on epistemic planning with propositional DEL (Section 9.1), results exist concerning the undecidability of several classes of epistemic planning problems, but decidability and complexity results also exist. It is clear that the negative results apply in the richer setting of this paper. In [55], we show that some of the positive decidability results can be established in the decidable finite-agent setting of dynamic term-modal logic (i.e., decidability for single-agent planning and multi-agent planning with non-modal preconditions). It is an open question whether any other decidability results can be extended as well, and the complexity of first-order epistemic planning has not been studied.
2. *Reasoning about schematic actions.* In extension to defining first-order variants of action models, it was natural to define action schemas to obtain succinct action representations. These action schemas are however not described by the dynamic languages and logics introduced. We find it an interesting question how the languages and logics should be altered to obtain a logic of action schemas. Constructing such a logic could possibly draw connections to recent work on *Arbitrary Public Announcement Logic* and its generalizations, cf. e.g. [5, 29].

⁷We thank a reviewer for pointing this out.

⁸We hold a reservation as satisfiability is not generally preserved by the translation. In the class of TML models with exactly 2 agents (characterized by axiom N for $n = 2$) wlog called α and β , with constants a and b locally rigid, but non-identical (characterized by $\exists x \exists y((x \neq y) \wedge (x = a) \wedge (y = b) \wedge \forall z K_z((x = a) \wedge (y = b)))$), and satisfying for $i, j \in \{\alpha, \beta\}, i \neq j$, $\forall x, y, z \in W$, if $xR_i y$ and $xR_j z$, then $yR_i z$ (characterized by $\forall x \forall y(((x \neq y) \wedge \hat{K}_x \top \wedge \hat{K}_y \varphi) \rightarrow K_x \hat{K}_y \varphi)$), the formula $\exists x \exists y((x \neq y) \wedge \hat{K}_y \top \wedge \hat{K}_x \hat{K}_x \top \wedge K_x K_x \neg \hat{K}_x \top)$ is satisfiable. However, the translation of the latter is not satisfiable in the class of FOML models characterized by the translation of the three former.

3. *Supporting other planning features.* A possibly fruitful avenue for future research is to devise a first-order *probabilistic* DEL framework for probabilistic epistemic planning. In the standard planning literature, probabilistic PDDL is often used to support probabilistic effects, allowing the specification of Markov decision processes [91]. There is a rich literature on probabilistic propositional DEL on which a first-order setting for probabilistic epistemic planning could be based (for an overview, see [8, Appendix L]). Other well-known planning features, such as numeric fluents, temporal aspects, etc., could also be integrated.

Finally, we briefly recall what we see as the main contributions of the paper:

1. *A first-order dynamic epistemic logic.* The paper develops novel dynamics for a variant of term-modal logic with the addition of first-order action models. It thereby generalizes propositional DEL to a setting allowing full first-order epistemic reasoning about both objects and agents.
2. *A compact epistemic domain definition language.* As the epistemic planning formalism developed builds on first-order logic, it allows for a compact specification of domain dynamics via *epistemic action schemas*. Such schematization is inspired by that used in PDDL, and to the best of our knowledge, it provides the most compact representation of actions available in the DEL framework. The setting conservatively extends propositional DEL, in the sense that it contains it as a special case, inheriting the ingredients of the DEL planning framework.
3. *Expressive, yet decidable axiom systems for reasoning about epistemic actions.* On the reasoning side, the paper develops static and dynamic axiom systems that are well-behaved. Although the logical languages proposed are fairly expressive, it is shown that sound, complete and decidable systems exist for several natural classes of models.

Acknowledgements

We sincerely thank the three anonymous reviewers for their insightful questions, comments and criticisms: We appreciate your time and efforts.

The Center for Information and Bubble Studies is funded by the Carlsberg Foundation. RKR was partially supported by the DFG-ANR joint project Collective Attitude Formation [RO 4548/8-1].

A Proof Appendix

A.1 Term-Modal Logic

This section establishes the results stated in Section 8.1. The logic K is well-behaved, with standard techniques for establishing strong completeness carrying over from the propositional and quantified modal logic cases. Therefore, the section presents only proof strategy, with non-standard elements given special attention. Full details may be found in [1].

The involved notions are standard (see e.g. [12, 18, 44]), but we remark that a formula φ is *valid* over a class of frames X iff for every frame $F = (D, W, R) \in X$, every interpretation I over F , every world $w \in W$ and every valuation v , it is the case that $M, w \models_v \varphi$. That φ is a *semantic consequence* of the formula-set Γ over a class X is written $\Gamma \models_X \varphi$. For φ provable from the assumptions Γ in the logic Λ , write $\Gamma \vdash_\Lambda \varphi$. In both cases, when $\Gamma = \emptyset$, it is omitted.

A.1.1 Soundness

Proposition 3. *The system K is sound with respect to the class F of all frames: for all $\varphi \in \mathcal{L}$, if $\vdash_K \varphi$, then $\models_F \varphi$.*

Proof. The proof is standard: the axioms of K are shown valid over F and the rules of inference are shown to preserve validity. To give a feel, arguments follow for the K axiom and the Barcan Formula.

K: Let M be a model based on an arbitrary frame $F \in F$, let $w \in M$ and let v be a valuation; let $K_t \top, \varphi, \psi \in \mathcal{L}$. To show that $M, w \models_v K_t(\varphi \rightarrow \psi) \rightarrow (K_t\varphi \rightarrow K_t\psi)$, assume $M, w \models_v K_t(\varphi \rightarrow \psi)$. As $K_t \top \in \mathcal{L}$, $\llbracket t \rrbracket_w^{L,v} \in D_{\text{agt}}$ by assumption. Hence F contains an accessibility relation $R_{\llbracket t \rrbracket_w^{L,v}}$. Having fixed the accessibility relation going through the term t to the agent domain, the argument is standard: By the semantics of K_t , $M, w' \models_v \varphi \rightarrow \psi$ for every $w' \in M$ with $w' \in R_{\llbracket t \rrbracket_w^{L,v}}(w)$. Hence $M, w' \models_v \neg\varphi$ or $M, w' \models_v \psi$. If all such w' satisfies φ , $M, w \models_v K_t\varphi$; but then each w' must also satisfy ψ , so $M, w \models_v K_t\psi$, and hence $M, w \models_v K_t\varphi \rightarrow K_t\psi$. Else, some such w' satisfies $\neg\varphi$; then $M, w \models_v \neg K_t\varphi$, so $M, w \models_v K_t\varphi \rightarrow K_t\psi$.

BF: Let M, w, v, φ and t be as above. Pick a variable $x \neq t$ and assume that $M, w \models_v \forall x K_t\varphi$. Then for all x -variants v' of v , $M, w \models_{v'} K_t\varphi$ (i.e., intuitively, if x is free in φ so that $K_t\varphi(x)$ defines a predicate, all elements in the $\tau(x)$ -domain of w fall in this predicate's extension). From $M, w \models_{v'} K_t\varphi$, it follows that for all $w' \in R_{\llbracket t \rrbracket_w^{L,v}}(w)$, $M, w' \models_{v'} \varphi$ (intuitively, as v' is an arbitrary x -variant v , all $\tau(x)$ -elements existing in w' fall in the extension of $\varphi(x)$). This would not hold if elements could exist in w' that do not exist in w . As v' is an arbitrary x -variant of v , it follows that $M, w' \models_{v'} \forall x \varphi$ (again, illegitimate if new elements could spring to existence). As w' was arbitrary from $R_{\llbracket t \rrbracket_w^{L,v}}(w)$, finally $M, w \models_v K_t\forall x \varphi$. \square

A.1.2 Completeness

This section establishes that the system K is *strongly complete* with respect to the class F of all frames. I.e.,

$$\text{for all } \Gamma \subseteq \mathcal{L}, \text{ for all } \varphi \in \mathcal{L}, \text{ if } \Gamma \models_F \varphi, \text{ then } \Gamma \vdash_K \varphi.$$

This follows as a corollary of the section's main result, the *Canonical Class Theorem* (Theorem 3) which states that any normal term-modal logic is strongly complete with respect to its *canonical class*.

The theorem is established by appeal to the following well-known⁹ proposition linking satisfaction and completeness:

Proposition 4. *A logic Λ is strongly complete with respect to a class of structure S iff every Λ -consistent set of formulas is satisfiable on some $s \in S$.*

By this proposition, a completeness proof can be undertaken as an existence proof: For a consistent set of formulas Γ , a satisfying model from the appropriate class must be found. In the propositional case, one model is constructed for all consistent sets simultaneously, giving rise to the propositional *Canonical Model Theorem* (see e.g. [12]): any normal propositional modal logic is strongly complete with respect to its *canonical model*.

The present proof cannot rely on single canonical model. As variables are semantically rigid and any signature Σ includes identity, the same identity statements between variables are true across all worlds of any model-valuation pair. A canonical model defined as usual would not satisfy this:

⁹See e.g. [12, p. 194].

with consistent sets forming the basis of worlds, if two worlds are disconnected by all accessibility relations, then they need not satisfy the same identity statements between variables. Hence, a rigid variable valuation cannot be defined. Further, different K -consistent sets may give rise to different domains. Hence, non-constant domains result, and the construction is thus not of the appropriate class. Therefore, our construction is of a canonical model *per* consistent set, resulting in a *canonical class*.

The construction contains first-order aspects irrelevant in the propositional case and term-modal logical aspects irrelevant to the standard quantified case, but the approach is familiar: worlds are maximally consistent sets that bear witnesses, ensured constructable by Lindenbaum-like lemmas; domains are equivalence classes of variables induced by identity statements; and canonical accessibility relations, interpretation and valuation are defined as expected. That the canonical accessibility relations are well-defined requires an additional lemma, but a familiar Existence Lemma facilitates a familiar Truth Lemma, which in combination with the above Proposition 4 yields the main result.

A.1.2.1 Canonical Worlds Fix a signature $\Sigma = (\mathcal{V}, \mathcal{C}, \mathcal{R}, \mathcal{F}, \mathfrak{t})$, its language \mathcal{L} and a normal term-modal logic $\Lambda \subseteq \mathcal{L}$. When a set $\Gamma \subseteq \mathcal{L}$ is *maximal Λ -consistent* (defined as usual [12]), call Γ a Λ -mcs.

Maximal consistency does not suffice for a set to be a canonical world in the first-order case. It must also be ensured that whenever a formula of the form $\neg\forall x\varphi$ is included in Γ , then Γ must bear witness to this “falsity” of $\forall x\varphi$:¹⁰

Definition 19. A set $\Gamma \subseteq \mathcal{L}$ *bears witnesses* if for every $\varphi \in \mathcal{L}$, for every variable x , there is some variable y such that $(\varphi(y/x) \rightarrow \forall x\varphi) \in \Gamma$.

If a set Γ bears witnesses, then so does every super-set of Γ . If Γ is a Λ -mcs that bears witnesses and contains $\neg\forall x\varphi$, then for some $y \in \mathcal{V}$, $\neg\varphi(y/x) \in \Gamma$.

To ensure that every Λ -mcs can be extended to one bearing witnesses, countably infinite sets of both agent and object variables beyond those in \mathcal{V} are needed. Define the extended signature Σ^+ as $(\mathcal{V}^+, \mathcal{C}, \mathcal{R}, \mathcal{F}, \mathfrak{t}^+)$ where $\mathcal{V} \subseteq \mathcal{V}^+$, $\mathfrak{t}^+(x) = \mathfrak{t}(x)$ for all $x \in \mathcal{V} \cup \mathcal{C} \cup \mathcal{R} \cup \mathcal{F}$ and both $(\mathfrak{t}^+)^{-1}(\text{agt}) \cap \mathcal{V}^+ \setminus \mathcal{V}$ and $(\mathfrak{t}^+)^{-1}(\text{obj}) \cap \mathcal{V}^+ \setminus \mathcal{V}$ are countably infinite. Let \mathcal{L}^+ be the term-modal language based on Σ^+ . Then $\mathcal{L} \subseteq \mathcal{L}^+$. The following two lemmas then ensure that the worlds of the canonical models are constructable:

Lemma 1 (Lindenbaum). *If $\Gamma \subseteq \mathcal{L}$ is Λ -consistent, then there is a Λ -mcs Γ' such that $\Gamma \subseteq \Gamma'$.*

Lemma 2 (Witnessed). *If $\Gamma \subseteq \mathcal{L}$ is Λ -consistent, then there is a set $\Gamma^+ \subseteq \mathcal{L}^+$ such that $\Gamma \subseteq \Gamma^+$ and Γ^+ bears witnesses.*

A.1.2.2 Canonical Models To avoid the issue remarked in this section’s introduction, a canonical model is defined per Λ -mcs, ensuring that all worlds share its *identity theory*:

Definition 20. The sets $\Gamma, \Gamma' \subseteq \mathcal{L}^+$ have the same *identity theory* if for all $x, y \in \mathcal{V}^+$, $(x = y) \in \Gamma$ iff $(x = y) \in \Gamma'$.

Definition 21. Let $\Lambda \subseteq \mathcal{L}$ be a normal term-modal logic. Let $\Gamma \subseteq \mathcal{L}$ be Λ -consistent and let $\Gamma^* \subseteq \mathcal{L}^+$ be maximal Λ -consistent, witness bearing and such that $\Gamma \subseteq \Gamma^*$ (existing by Lemmas 1 and 2). The *canonical model* for (Λ, Γ^*) is $M_{(\Lambda, \Gamma^*)} = (D, W, R, I)$ such that

¹⁰Witnesses bearing is called the \forall -property in [44, p. 257]; that the set is *saturated* is also used in the literature.

1. $D := D_{\text{agt}} \dot{\cup} D_{\text{obj}} := \{[x] : x \in (\mathfrak{t}^+)^{-1}(\text{agt}) \cap V^+\} \dot{\cup} \{[y] : y \in (\mathfrak{t}^+)^{-1}(\text{agt}) \cap V^+\}$ where $[z] := \{z' \in V^+ : (z = z') \in \Gamma^*\}$.
2. W is the set of all maximal Λ -consistent, witness bearing sets of formulas from \mathcal{L}^+ that share identity theory with Γ^* .
3. $R : D_{\text{agt}} \rightarrow \mathcal{P}(W \times W)$ such that for all $\alpha \in D_{\text{agt}}$, $(w, w') \in R(\alpha)$ iff for every formula $K_x \varphi \in \mathcal{L}^+$ with $x \in \alpha$, if $K_x \varphi \in w$, then $\varphi \in w'$,
4. and
 - (a) $I(r, w) = \{([x_1], \dots, [x_n]) \in \prod_{i=1}^{\text{len}(\mathfrak{t}(r))} D_{\mathfrak{t}_i(r)} : r(x_1, \dots, x_n) \in w\}$, for all $r \in \mathbb{R}$;
 - (b) $I(f, w) = \{([x_1], \dots, [x_n]) \in \prod_{i=1}^{\text{len}(\mathfrak{t}(f))} D_{\mathfrak{t}_i(f)} : (f(x_1, \dots, x_{n-1}) = x_n) \in w\}$, for all $f \in \mathbb{F}$;
 - (c) $I(c, w) = \{([x]) \in D_{\mathfrak{t}(c)} : (c = x) \in w\}$, for all $c \in \mathbb{C}$.

The *canonical valuation* v for (Λ, Γ^*) is given by $v(x) = [x]$ for all $x \in V^+$.

A.1.2.3 Lemmas: Uniformity, Existence and Truth The canonical model for (Λ, Γ^*) is a model for \mathcal{L} . Notably, the domain is well-defined by the identity theory sharing requirement and a two-partition by the inclusion of the DD axiom. Further, $I(c, w)$ is well-defined as for every world w , there exists some $x \in V^+$ for which $(c = x) \in w$. See [1] for details. Foremost, the map R is well-defined, as is ensured by the following lemma:

Lemma 3 (Uniformity). *Let $K_x \varphi \in w \in W$ with $v(x) = \alpha$. Then for all $y \in V^+$ for which $v(x) = v(y)$, also $K_y \varphi \in w$.*

Proof. Assume $K_x \varphi \in w \in W$ with $v(x) = \alpha$, and let $v(x) = v(y)$. Then $[x] = [y]$, so by identity theory sharing assumption, $(x = y) \in w'$ for every $w' \in W$; in particular, $(x = y) \in w$. By PS, $(x = y) \rightarrow (K_x \varphi \rightarrow K_y \varphi) \in w$. By MP, $(K_x \varphi \rightarrow K_y \varphi) \in w$ and by MP again, $K_y \varphi \in w$. \square

As in the propositional case, the proof of the Truth Lemma below relies on the below Existence Lemma. A proof for standard first-order modal logic may be found in [44]; details for term-modal logic may be found in [1].

Lemma 4 (Existence). *If $w \in W$ and $\neg K_x \varphi \in w$, then there exists a $w' \in W$ such that $(w, w') \in R_{\llbracket x \rrbracket_w^{L,v}}$ and $\varphi \in w'$.*

Lemma 5 (Truth). *For all $\varphi \in \mathcal{L}^+$, for all $w \in W$, and for the canonical v , $M_{(\Lambda, \Gamma^*)}, w \models_v \varphi$ iff $\varphi \in w$.*

Proof. The proof proceeds by induction on the complexity of φ . For the quantified formulas, appeal is made to w bearing witnesses. The negated modal case relies on the Existence Lemma. See [1] for full details. \square

A.1.2.4 Canonical Class Theorem The canonical models defined facilitate the application of Proposition 4 to conclude strong completeness of Λ with respect to its *canonical class*:

Definition 22. The *canonical class* of models for the normal term-modal logic Λ is the set C_Λ of canonical models $M_{(\Lambda, \Gamma^*)}$ for Λ -consistent $\Gamma \subseteq \mathcal{L}$.

Theorem 3 (Canonical Class Theorem). *Any normal term-modal logic Λ is strongly complete with respect to its canonical class.*

Proof. By Proposition 4, it suffices to find for each Λ -consistent set Γ some $s \in C_\Lambda$ that satisfies Γ . One such is $(M_{(\Lambda, \Gamma^*)}, \Gamma^*)$, which exists by the Lindenbaum and Witnessed Lemmas. As $\Gamma \subseteq \Gamma^*$, the Truth Lemma ensure that $(M_{(\Lambda, \Gamma^*)}, \Gamma^*) \models_v \Gamma$ for v the canonical valuation. \square

Corollary 1 (Completeness). *The logic K is strongly complete with respect to the class of all frames F .*

Proof. A frame $F \in \mathcal{F}$ that satisfies the K -consistent set Γ is the frame of the canonical model $M_{(K, \Gamma^*)}$: Γ is satisfied at Γ^* under the canonical valuation. \square

A.1.3 Frame Characterization Proofs

For illustrative purposes, we show two of the claims made in Table 2, Section 8.1.4.

Proposition 5. *For $\varphi \in \mathcal{L}$, $\forall x(K_x\varphi \rightarrow K_xK_x\varphi)$ is valid on the frame $F = (D, W, R)$ if, and only if, $R(\alpha)$ is transitive for every $\alpha \in D_{\text{agt}}$.*

Proof. \Leftarrow : Let M be build on the frame F in which R_α is transitive for all $\alpha \in D_{\text{agt}}$. Let v be an arbitrary valuation and assume $M, w \models_v K_x\varphi$. Then $M, w' \models_v \varphi$ for all $w' \in R_{v(x)}(w)$. For a contradiction, assume $M, w \models_v \neg K_xK_x\varphi$. Then there exists a $w^* \in R_{v(x)}(w)$ such that $M, w^* \models_v \neg K_x\varphi$, and hence there exists a $w^{**} \in R_{v(x)}(w^*)$ such that $M, w^{**} \models_v \neg\varphi$. But $R_{v(x)}$ is transitive, so $w^{**} \in R_{v(x)}(w)$. Hence w^{**} satisfies both φ and $\neg\varphi$. On pain of contradiction, $M, w \models_v K_xK_x\varphi$. As v was arbitrary, $M, w \models_v \forall x(K_x\varphi \rightarrow K_xK_x\varphi)$. \Rightarrow : By contraposition. \square

Proposition 6. *The formula $\exists x_1, \dots, x_n \left((\bigwedge_{i \leq n} K_{x_i} \top) \wedge \left(\bigwedge_{i, j \leq n, i \neq j} x_i \neq x_j \right) \wedge \forall y (K_y \top \rightarrow \bigvee_{i \leq n} y = x_i) \right)$ is valid on the frame $F = (W, D, R)$ if, and only if, $|D_{\text{agt}}| = n$.*

Proof. Notice first that the formula, call it N , is only well-formed iff the variables x_1, \dots, x_n, y are all of type agt , ensured by them appearing as modal operator subscripts. This ensures that the quantifications range only over D_{agt} .

\Leftarrow : Assume given a pointed model (M, w) build on a frame $F = (W, D, R)$ with $|D_{\text{agt}}| = n$. Assume D_{agt} enumerated such that $D_{\text{agt}} = \{\alpha_1, \dots, \alpha_n\}$. Let v be an arbitrary valuation. We argue that $M, w \models_v N$. Let v' be the valuation identical to v on all points, except for each $i \leq n$, $v'(x_i) = \alpha_i$. Then $M, w \models_{v'} \left((\bigwedge_{i \leq n} K_{x_i} \top) \wedge \left(\bigwedge_{i, j \leq n, i \neq j} x_i \neq x_j \right) \wedge \forall y (K_y \top \rightarrow \bigvee_{i \leq n} y = x_i) \right)$, as it satisfies each conjunct: First, $M, w \models_{v'} \bigwedge_{i \leq n} K_{x_i} \top$, trivially. Second, $M, w \models_{v'} \bigwedge_{i, j \leq n, i \neq j} x_i \neq x_j$ as $v'(x_i) \neq v'(x_j)$ for all $i, j \leq n$, $i \neq j$, by construction of v' . Third and finally, $M, w \models_{v'} \forall y (K_y \top \rightarrow \bigvee_{i \leq n} y = x_i)$: as y is of type agt , for any y -variant v'' of v' , $v'' \in D_{\text{agt}}$, but then $v''(y) = v''(x_i)$ for some $i \leq n$, by construction of v' , satisfying the antecedent.

\Rightarrow : Assume given a pointed model (M, w) build on a frame $F = (W, D, R)$ with $|D_{\text{agt}}| \neq n$. Let v be an arbitrary valuation. We argue that not $M, w \models_v N$, as (M, w) will falsify either the second

or the third conjunct (the first conjunct is satisfied: as each variable x_i is of type agt for all $i \leq n$, each $K_{x_i} \top$ is satisfied at w under any valuation). If $|D_{\text{agt}}| < n$, then under any valuation v , (M, w) will falsify the second conjunct: as each variable x_i is of type agt , $v(x_i) \in D_{\text{agt}}$ for all $i \in 1, \dots, n$. But then $v(x_i) = v(x_j)$ for at least two $i, j \leq n, i \neq j$. But then $M, w \models_v x_i = x_j$, contrary to the second conjunct. If $|D_{\text{agt}}| > n$, then under any valuation v , (M, w) will falsify the third conjunct, as there exists a y -variant v' of v such that $v'(y) \neq v(x_i)$ for any $i \leq n$. The existence of this y -variant v' is ensured by $|D_{\text{agt}}| > n$, which implies that $D_{\text{agt}}/v(x_i): i \leq n \neq \emptyset$, so that we can assume $v'y \in D_{\text{agt}}/v(x_i): i \leq n$. \square

A.1.4 Decidability

Proposition 7. *Let $K_{n/m}$, K_n and K be given in \mathcal{L} , based on the signature Σ . Let $\mathcal{L}_{\text{agt}} \subseteq \mathcal{L}$ contain all formulas containing only agent-terms, $t \in \mathfrak{t}^{-1}(\text{agt})$.*

1. *For all $\varphi \in \mathcal{L}$, it is decidable whether $\vdash_{K_{n/m}} \varphi$ or not.*
2. *a) For all $\varphi \in \mathcal{L}_{\text{agt}}$, it is decidable whether $\vdash_{K_n} \varphi$ or not. b) In general, $\vdash_{K_n} \varphi$ is undecidable.*
3. *In general, $\vdash_K \varphi$ is undecidable.*

Proof. **1.** $K_{n/m}$ is sound and complete w.r.t. $F_{n/m}$. To check the validity of any $\varphi \in \mathcal{L}$ over $F_{n/m}$ is a finite procedure: Up to isomorphism, all $F \in F_{n/m}$ share domain $D = D_{\text{agt}} \cup D_{\text{obj}}$, $|D_{\text{agt}}| = n$, $|D_{\text{obj}}| = (m - n)$. There are finitely many non-logical symbols in φ ; symbols not in φ are irrelevant to its satisfaction. With D fixed, any $w \in F$ will be assigned one of finitely many extensions of φ 's non-logical symbols: thus, the maximal set of distinct φ -relevant worlds W_φ is finite. As φ has modal depth k , whether $M, w \models_v \varphi$ depends on at most all worlds within k steps from w . Checking whether $M, w \models_v \varphi$ is thus a finite procedure for all formulas given the finiteness of D . Finally, up to bisimulation, the set of graphs over W_φ and $\{R(\alpha), \alpha \in D_{\text{agt}}\}$ with maximal path length k is finite: hence, the set of needed to be checked pointed models is finite. **2a.** For any $\varphi \in \mathcal{L}_{\text{agt}}$, φ is a theorem of K_n iff it is a theorem of $K_{n/m}$, for any $m > n$. For such φ , to determine whether $\vdash_{K_n} \varphi$, we can thus check whether $\vdash_{K_{n/n+1}} \varphi$, which is decidable by 1. **2b** and **3.** General undecidability for K_n and K follows as both contain unrestricted first-order logic for the arbitrary object domain. \square

A.2 Dynamic Term-Modal Logic: Completeness through Translation

The completeness proof for the dynamic logic $K + \text{AM}$ is based on a reduction argument. The argument relies on the existence of so-called reduction axioms for the dynamic language \mathcal{L}_{AM} . The axioms used for this specific proof are listed in Table 3 and can be used to translate every formula from the dynamic language \mathcal{L}_{AM} into a provably equivalent \mathcal{L} -formula. Given this translation, the completeness of the dynamic logic follows from the known completeness of the static logic K , established in Corollary 1. The building blocks of the specific reduction argument required to prove completeness for $K + \text{AM}$ are provided below.

First, we provide a translation that by finite iterative application to any formula in the dynamic language \mathcal{L}_{AM} results in a formula from the static language \mathcal{L} . The translation is left-to-right: a formula occurring on the left is translated to the formula on the right.

Definition 23. The *translation* $\tau : \mathcal{L}_{AM} \rightarrow \mathcal{L}_{AM}$ is defined as follows:

$$\begin{aligned}
\tau((t_1 = t_2)) &= (t_1 = t_2) \\
\tau(r(t_1, \dots, t_n)) &= r(t_1, \dots, t_n) \\
\tau(\neg\varphi) &= \neg\tau(\varphi) \\
\tau(\varphi \wedge \psi) &= \tau(\varphi) \wedge \tau(\psi) \\
\tau(K_t\varphi) &= K_t\tau(\varphi) \\
\tau(\forall x\varphi) &= \forall x\tau(\varphi) \\
\tau([A, e]r(t_1, \dots, t_n)) &= \tau(\text{pre}(e) \rightarrow \text{post}^A(e)(r(t_1, \dots, t_n))) \\
\tau([A, e]\neg\varphi) &= \tau(\text{pre}(e) \rightarrow \neg[A, e]\varphi) \\
\tau([A, e](\varphi \wedge \psi)) &= \tau([A, e]\varphi \wedge [A, e]\psi) \\
\tau([A, e]K_t\varphi) &= \tau(\text{pre}(e) \rightarrow \bigwedge_{e' \in E^A} (Q(e, e')[x^* \mapsto t]K_t[A, e']\varphi)) \\
\tau([A, e]\forall x\varphi) &= \tau(\text{pre}(e) \rightarrow \forall x[A, e]\varphi) \\
\tau([A, e][A', e']\varphi) &= \tau([A, e \circ A', e']\varphi)
\end{aligned}$$

Next, we adapt the formula complexity function introduced by [28].

Definition 24. The *complexity* $c : \mathcal{L}_{AM} \rightarrow \mathbb{N}$ is defined as follows, where $\text{GA}(\mathcal{L})$ abbreviates $\text{GroundAtoms}(\mathcal{L})$:

$$\begin{aligned}
c(r(t_1, \dots, t_n)) &= 1 \\
c(\neg\varphi) &= 1 + c(\varphi) \\
c(\varphi \wedge \varphi') &= 1 + \max(c(\varphi), c(\varphi')) \\
c(K_t\varphi) &= 1 + c(\varphi) \\
c(\forall x\varphi) &= 1 + c(\varphi) \\
c([A, e]\varphi) &= (4 + c(A)) \cdot c(\varphi) \\
c(A) &= \max \left(\bigcup_{e, e' \in E, r(t_1, \dots, t_n) \in \text{GA}(\mathcal{L})} \{c(\text{pre}^A(e))\} \cup \{c(\text{post}^A(e)(r(t_1, \dots, t_n)))\} \cup \{c(Q(e, e'))\} \right)
\end{aligned}$$

A standard ordering lemma ensures that the right side of a given reduction axiom is indeed less complex than the left side.

Lemma 6. For all φ, ψ and χ :

1. $c(\psi) \geq c(\varphi)$ if $\varphi \in \text{Sub}(\psi)$ (where $\text{Sub}(\psi)$ is the set of subformulas of ψ)
2. $c([A, e]r(t_1, \dots, t_n)) > c(\text{pre}(e) \rightarrow \text{post}(e)(r(t_1, \dots, t_n)))$
3. $c([A, e]\neg\varphi) > c(\text{pre}(e) \rightarrow \neg[A, e]\varphi)$
4. $c([A, e](\varphi \wedge \psi)) > c([A, e]\varphi \wedge [A, e]\psi)$
5. $c([A, e]K_t\varphi) > c(\text{pre}(e) \rightarrow \bigwedge_{e' \in E} (Q(e, e')[x^* \mapsto t]K_t[A, e']\varphi))$

6. $c([A, e] \forall x \varphi) > c(\text{pre}(e) \rightarrow \forall x[A, e] \varphi)$
7. $c([A, e][A', e'] \varphi) > c([A, e \circ A', e'] \varphi)$

Proof. The proofs are straightforward, along the lines of those provided in [28, Chapter 7]. \square

The complexity function c induces an ordering of \mathcal{L}_{AM} formulas which is used to prove the following Lemma, stating that the two sides of a reduction axiom are indeed provably equivalent.

Lemma 7. *For all $\varphi \in \mathcal{L}_{AM}$: $\vdash_{K+AM} \varphi \leftrightarrow \tau(\varphi)$.*

Proof. The proof is by induction on the complexity $c(\varphi)$. It is similar to the one provided in [28, Chapter 7]. \square

The completeness of $K + AM$ (Corollary 2) follows from the soundness of the dynamic proof system, Lemma 7 and the completeness of the static sub-system (Corollary 1). The argument, which is standard, is as follows.

Proposition 8. $\models \varphi$ implies $\vdash_{K+AM} \varphi$, for all $\varphi \in \mathcal{L}_{AM}$.

Proof. Suppose $\models \varphi$. Since $\vdash_{K+AM} \varphi \leftrightarrow \tau(\varphi)$ (Lemma 7), we have $\models \varphi \leftrightarrow \tau(\varphi)$ by the soundness of the proof system $K + AM$. Thus $\models \tau(\varphi)$. The formula $\tau(\varphi)$ does not contain any action model modalities. Given $\models \tau(\varphi)$, by the completeness of K (Corollary 1), it follows that $\vdash_K \tau(\varphi)$. As K is a subsystem of $K + AM$, we thus have $\vdash_{K+AM} \tau(\varphi)$. Since $\vdash_{K+AM} \varphi \leftrightarrow \tau(\varphi)$ and $\vdash_{K+AM} \tau(\varphi)$, it follows that $\vdash_{K+AM} \varphi$. \square

The completeness result for any system extending $K + AM$ with frame-characterizing axioms follows from the same type of argument.

References

- [1] A. Achen. Putting the Agents Back in the Domain: A Two-Sorted Term-Modal Logic. Bachelor thesis, University of Copenhagen, May 2017.
- [2] M. B. Andersen, T. Bolander, and M. H. Jensen. Conditional Epistemic Planning. In *Logics in Artificial Intelligence*, pages 94–106. Springer, 2012.
- [3] G. Aucher. An Internal Version of Epistemic Logic. *Studia Logica*, 94(1):1–22, 2010.
- [4] G. Aucher and T. Bolander. Undecidability in Epistemic Planning. In *IJCAI-13*, pages 27–33, 2013.
- [5] P. Balbiani, A. Baltag, H. v. Ditmarsch, A. Herzig, T. Hoshi, and T. de Lima. What Can We Achieve by Arbitrary Announcements?: A Dynamic Take on Fitch’s Knowability. In *TARK ’07: Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 42–51. ACM, 2007.
- [6] A. Baltag and L. S. Moss. Logics for Epistemic Programs. *Synthese*, 139(2):165–224, 2004.
- [7] A. Baltag, L. S. Moss, and S. S. Solecki. The Logic of Public Announcements, Common Knowledge, and Private Suspicions (extended abstract). In *TARK ’98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 43–56. Morgan Kaufmann Publishers, 1998.

- [8] A. Baltag and B. Renne. Dynamic Epistemic Logic. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition, 2016.
- [9] C. Baral, T. Bolander, H. van Ditmarsch, and S. McIlrath. Epistemic Planning (Dagstuhl Seminar 17231). *Dagstuhl Reports*, 7(6):1–47, 2017.
- [10] J. van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.
- [11] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [12] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [13] T. Bolander. Seeing is Believing: Formalising False-Belief Tasks in Dynamic Epistemic Logic. In A. Herzig and E. Lorini, editors, *European Conference on Social Intelligence (ECSI 2014)*, volume 1283, pages 87–107. CEUR Workshop Proceedings, vol. 1283, 2014.
- [14] T. Bolander. A Gentle Introduction to Epistemic Planning: The DEL Approach. In *9th Workshop on Methods for Modalities*, volume 243, pages 1–22, 2017.
- [15] T. Bolander and M. B. Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.
- [16] T. Bolander and M. B. Andersen. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.
- [17] T. Bolander, M. H. Jensen, and F. Schwarzenrüber. Complexity Results in Epistemic Planning. In *IJCAI-15*, pages 2791–2797, 2015.
- [18] T. Brauner and S. Ghilardi. First-order Modal Logic. In F. W. P. Blackburn, J. van Benthem, editor, *Handbook of Modal Logic*. Elsevier Science, 2007.
- [19] M. Brenner. A Multiagent Planning Language. In *Proc. of the Workshop on PDDL, ICAPS*, volume 3, pages 33–38, 2003.
- [20] L. Carlson. Quantified Hintikka-style Epistemic Logic. *Synthese*, 74:223–262, 1988.
- [21] T. Charrier, B. Maubert, and F. Schwarzenrüber. On the Impact of Modal Depth in Epistemic Planning. In *IJCAI-16*, pages 1030–1036, 2016.
- [22] T. Charrier and F. Schwarzenrüber. A Succinct Language for Dynamic Epistemic Logic. In *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems*, pages 123–131. International Foundation for Autonomous Agents and Multiagent Systems, 2017.
- [23] M. C. Cooper, A. Herzig, F. Maffre, F. Maris, and P. Régnier. A Simple Account of Multi-Agent Epistemic Planning. In *ECAI*, pages 193–201, 2016.
- [24] G. Corsi and E. Orlandelli. Free Quantified Epistemic Logics. *Studia Logica*, 101(6):1159–1183, Dec 2013.

- [25] G. Corsi and E. Orlandelli. Sequent Calculi for Indexed Epistemic Logics. In *ARQNL 2016: Automated Reasoning in Quantified Non-Classical Logics*, volume 1770 of *CEUR-WS*, pages 21–35, 2016.
- [26] G. Corsi and G. Tassi. A New Approach to Epistemic Logic. In E. Weber, D. Wouters, and J. Meheus, editors, *Logic, Reasoning, and Rationality*, pages 27–44, Dordrecht, 2014. Springer Netherlands.
- [27] H. van Ditmarsch, W. van der Hoek, and B. Kooi. Dynamic epistemic logic with assignment. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 141–148. ACM, 2005.
- [28] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [29] H. van Ditmarsch, W. van der Hoek, and L. B. Kuijer. Fully Arbitrary Public Announcements. In L. Beklemishev, S. Demri, and A. Máté, editors, *Advances in Modal Logic*, volume 11, pages 252–267, 2018.
- [30] H. van Ditmarsch and B. Kooi. Semantic Results for Ontic and Epistemic Change. In G. Bonanno, W. van der Hoek, and M. Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games, Vol. 3, pages 87–117. Amsterdam University Press, 2008.
- [31] T. Engesser, R. Mattmüller, B. Nebel, and M. Thielscher. Game Description Language and Dynamic Epistemic Logic Compared. In *IJCAI-18. International Joint Conferences on Artificial Intelligence Organization*, 2018.
- [32] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.
- [33] M. Fitting and R. L. Mendelsohn. *First-Order Modal Logic*. Springer, 1999.
- [34] M. Fitting, L. Thalmann, and A. Voronkov. Term-Modal Logics. *Studia Logica*, 69:133–169, 2001.
- [35] L. T. F. Gamut. *Logic, Language and Meaning. Volume 2: Intensional Logic and Logical Grammar*. The University of Chicago Press, 1996.
- [36] M. Ghallab, D. Nau, and P. Traverso. *Automated Planning: Theory and Practice*. Elsevier, 2004.
- [37] A. J. Grove. Naming and identity in epistemic logic Part II: a first-order logic for naming. *Artificial Intelligence*, 74(2):311 – 350, 1995.
- [38] A. J. Grove and J. Y. Halpern. Naming and Identity in Epistemic Logics Part I: The Propositional Case. *Journal of Logic and Computation*, 3(4):345–378, 1993.
- [39] M. Helmert. The Fast Downward Planning System. *Journal of Artificial Intelligence Research*, 26:191–246, 2006.
- [40] J. Hintikka. *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. College Publications, 2nd, 2005 edition, 1962.

- [41] J. Hintikka. Semantics for Propositional Attitudes. In J. W. Davis, D. J. Hockney, and W. K. Wilson, editors, *Philosophical Logic*, Synthese Library, pages 21–45. Springer, 1969.
- [42] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 3*, pages 1167–1174. ACM, 2002.
- [43] X. Huang, B. Fang, H. Wan, and Y. Liu. A General Multi-agent Epistemic Planner Based on Higher-order Belief Change. In *IJCAI-17*, pages 1093–1101, 2017.
- [44] G. Hughes and M. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.
- [45] W. Jamroga and T. Ågotnes. Constructive knowledge: What agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.
- [46] F. Kominis and H. Geffner. Beliefs In Multiagent Planning: From One Agent to Many. In *ICAPS*, pages 147–155, 2015.
- [47] F. Kominis and H. Geffner. Multiagent Online Planning with Nested Beliefs and Dialogue. In *Proc. ICAPS*, 2017.
- [48] B. Kooi. Dynamic term-modal logic. In J. van Benthem, S. Ju, and F. Veltman, editors, *A meeting of the minds, Proceedings of the Workshop on Logic, Rationality and Interaction, Beijing, 2007*, Texts in Computer Science 8, pages 173–185. Texts in Computer Science 8, College Publications, London, 2007.
- [49] D. L. Kovács. A Multi-Agent Extension of PDDL3.1. In *Proceedings of the 3rd Workshop on the International Planning Competition*, pages 19–27, 2012.
- [50] S. A. Kripke. The Undecidability of Monadic Modal Quantification Theory. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 8:113–116, 1962.
- [51] T. Le, F. Fabiano, T. C. Son, and E. Pontelli. EFP and PG-EFP: Epistemic Forward Search Planners in Multi-Agent Domains. In *ICAPS*, pages 161–170, 2018.
- [52] S. Lê Cong, S. Pinchinat, and F. Schwarzentruber. Small Undecidable Problems in Epistemic Planning. In *IJCAI-18*, pages 4780–4786, 2018.
- [53] H. S. van Lee, R. K. Rendsvig, and S. van Wijk. Intensional Protocols for Dynamic Epistemic Logic. *Journal of Philosophical Logic*, 48(6):1077–1118, Dec 2019.
- [54] A. O. Liberman and R. K. Rendsvig. Dynamic Term-Modal Logic for Epistemic Social Network Dynamics. In P. Blackburn, E. Lorini, and M. Guo, editors, *Logic, Rationality, and Interaction (LORI 2019)*, volume 11813 of *Lecture Notes in Computer Science*, pages 168–182. Springer, 2019.
- [55] A. O. Liberman and R. K. Rendsvig. Decidability Results in First-Order Epistemic Planning. In *IJCAI-20*, page to appear, 2020.
- [56] A. Lomuscio. QLb: una logica predicativa delle credenze. sintassi, semantica e una procedura di prova automatica dei teoremi. Master’s thesis, Politecnico di Milano, December 1995.

- [57] A. Lomuscio and M. Colombetti. QLB: A Quantified Logic for Belief. In J. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III. Agent Theories, Architectures, and Languages*, volume 1193 of *LNAI*, pages 71–85. Springer, 1997.
- [58] B. Löwe, E. Pacuit, and A. Witzel. DEL Planning and Some Tractable Cases. In *International Workshop on Logic, Rationality and Interaction*, pages 179–192. Springer, 2011.
- [59] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins. Pddl-the planning domain definition language, 1998.
- [60] L. S. Moss. Dynamic Epistemic Logic. In H. van Ditmarsch, J. Y. Halpern, W. van der Hoek, and B. Kooi, editors, *Handbook of Epistemic Logic*. College Publications, 2015.
- [61] C. J. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg. Planning Over Multi-Agent Epistemic States: A Classical Planning Approach. In *AAAI*, pages 3327–3334, 2015.
- [62] P. Naumov and J. Tao. Everyone Knows that Someone Knows: Quantifiers over Epistemic Agents. *Review of Symbolic Logic*, forthcoming, 2018.
- [63] E. Orlandelli and G. Corsi. Decidable Term-Modal Logics. In F. Belardinelli and E. Argente, editors, *Multi-Agent Systems and Agreement Technologies*, pages 147–162, Cham, 2018. Springer International Publishing.
- [64] A. Padmanabha. *Propositional Term Modal Logic*. PhD thesis, Homi Bhabha National Institute, 2019.
- [65] A. Padmanabha and R. Ramanujam. Model checking a logic over systems with regular sets of processes. In *Joint Proceedings of the 3rd Modelling Symposium (ModSym), Developmental Aspects of Intelligent Adaptive Systems (DIAS), and Educational Data Mining Practices in Indian Academia (EDUDM) co-located with 10th Innovations in Software Engineering (ISEC 2017)*, volume 1819 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2017.
- [66] A. Padmanabha and R. Ramanujam. Propositional Modal Logic with Implicit Modal Quantification. In M. A. Khan and A. Manuel, editors, *Logic and Its Applications (ICLA 2019)*, volume 11600 of *Lecture Notes in Computer Science*, pages 6–17. Springer Berlin Heidelberg, 2019.
- [67] A. Padmanabha and R. Ramanujam. The Monodic Fragment of Propositional Term Modal Logic. *Studia Logica*, 107(3):533–557, 2019.
- [68] A. Padmanabha and R. Ramanujam. Two variable fragment of Term Modal Logic. In P. Rossmanith, P. Heggernes, and J.-P. Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019)*, volume 138 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 30:1–30:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [69] A. Padmanabha and R. Ramanujam. Verifying Implicitly Quantified Modal Logic over Dynamic Networks of Processes. In D. V. Hung and M. D’Souza, editors, *Distributed Computing and Internet Technology (ICDCIT 2020)*, volume 11969 of *Lecture Notes in Computer Science*, pages 165–176. Springer, 2020.

- [70] E. P. Pednault. Formulating multiagent, dynamic-world problems in the classical planning framework. In *Reasoning about actions & plans*, pages 47–82. Elsevier, 1987.
- [71] R. P. Petrick and F. Bacchus. Extending the Knowledge-Based Approach to Planning with Incomplete Information and Sensing. In *ICAPS*, pages 2–11, 2004.
- [72] J. A. Plaza. Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z. W. Ras, editors, *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, pages 201–216, 1989.
- [73] R. Pliuškevičius and A. Pliuškevičienė. Decision Procedure for a Fragment of Mutual Belief Logic with Quantified Agent Variable. In F. Toni and P. Torrini, editors, *CLIMA VI, LNAI 3900*, pages 112–128. Springer-Verlag, 2006.
- [74] R. K. Rendsvig. Epistemic Term-Modal Logic. In M. Slavkovik, editor, *Proceedings of the 15th Student Session of The European Summer School in Logic, Language and Information*, pages 37–46, 2010.
- [75] R. K. Rendsvig. Towards a Theory of Semantic Competence. Master’s thesis, Dept. of Philosophy and Science Studies and Dept. of Mathematics, Roskilde University, 2011.
- [76] R. K. Rendsvig. Modeling Semantic Competence: A Critical Review of Frege’s Puzzle about Identity. In D. Lassiter and M. Slavkovik, editors, *New Directions in Logic, Language and Computation*, volume 7415 of *Lecture Notes in Computer Science*, pages 140–157. Springer, 2012.
- [77] R. K. Rendsvig. Pluralistic Ignorance in the Bystander Effect: Informational Dynamics of Unresponsive Witnesses in Situations calling for Intervention. *Synthese*, 191:2471–2498, 2014.
- [78] R. K. Rendsvig. Model Transformers for Dynamical Systems of Dynamic Epistemic Logic. In W. van der Hoek, W. H. Holliday, and W. F. Wang, editors, *Logic, Rationality, and Interaction (LORI 2015, Taipei)*, LNCS, pages 316–327. Springer, 2015.
- [79] R. K. Rendsvig. *Logical Dynamics and Dynamical Systems*. PhD thesis, Lund University, 2018.
- [80] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education Limited, 2016.
- [81] T. Sawasaki, K. Sano, and T. Yamada. Term-Sequence-Modal Logics. In P. Blackburn, E. Lorini, and M. Guo, editors, *Logic, Rationality, and Interaction (LORI 2019)*, volume 11813 of *Lecture Notes in Computer Science*, pages 244–258. Springer Berlin Heidelberg, 2019.
- [82] I. Sedlar. Term-Modal Logics of Evidence. In *Epistemic Logic for Individual, Social, and Interactive Epistemology (ESSLLI 2014)*, Tübingen, Germany, 2014.
- [83] G. Shtakser. Propositional Epistemic Logics with Quantification Over Agents of Knowledge. *Studia Logica*, 106(2):311–344, Apr 2018.
- [84] R. C. Sleight. Restricted Range in Epistemic Logic. *Journal of Philosophy*, 69:67–77, 1972.
- [85] L. Thalmann. *Term-Modal Logic and Quantifier-free Dynamic Assignment Logic*. PhD thesis, Uppsala University, 2000.

- [86] M. Thielscher. Gdl-iii: A description language for epistemic general game playing. In *IJCAI-17*, pages 1276–1282, 2017.
- [87] J. van Benthem, J. van Eijck, M. Gattinger, and K. Su. Symbolic model checking for dynamic epistemic logic - s5 and beyond. *Journal of Logic and Computation*, 28(2):367–402, 2018.
- [88] G. H. von Wright. *An Essay in Modal Logic*. North-Holland Publishing Company, 1951.
- [89] Y. Wang and J. Seligman. When Names Are Not Commonly Known: Epistemic Logic with Assignments. In *Advances in Modal Logic 12, proceedings of the 12th conference on "Advances in Modal Logic," held in Bern, Switzerland, August 27-31, 2018*, pages 611–628, 2018.
- [90] D. S. Weld. An Introduction to Least Commitment Planning. *AI Magazine*, 15(4):27, 1994.
- [91] H. L. Younes and M. L. Littman. PPDDL1.0: An Extension to PDDL for Expressing Planning Domains with Probabilistic Effects. *Techn. Rep. CMU-CS-04-162*, 2004.
- [92] Q. Yu, X. Wen, and Y. Liu. Multi-Agent Epistemic Explanatory Diagnosis via Reasoning about Actions. In *IJCAI-13*, pages 1183–1190, 2013.