

Decidable Multi-agent Epistemic Planning: A Situation Calculus Approach

Qihui Feng, Gerhard Lakemeyer

Knowledge-Based Systems Group, RWTH Aachen University, Germany
{feng, gerhard}@kbsg.rwth-aachen.de

Abstract

Multi-agent epistemic planning (MEP) is the task of generating action sequences that achieve goals specified over both the physical world and agents' mental states. It plays an important role in research domains such as game theory, computational economics, and cognitive science. While dynamic epistemic logic (DEL) provides an expressive framework for MEP, it requires complete, model-based specifications of the initial state and action effects, and suffers from undecidability due to the unbounded nesting of beliefs.

In this work, we propose a modal variant of the situation calculus that captures much of the expressive power of the DEL approach. Inspired by the cognitive concept Theory of Mind (ToM), we introduce action theories with hierarchical structures, allowing agents to reason about other agents' action theories up to bounded depths. We develop a regression method that reduces reasoning about future states to reasoning about the initial state. By preserving bounded-order ToM throughout the regression process, our approach ensures the decidability of the planning problem. Finally, we propose an algorithm to find the optimal solution, namely, to find the shortest action sequence that achieves the goal.

1 Introduction

Multi-agent epistemic planning (MEP) refers to the task of generating action sequences to achieve goals formulated not only in terms of the world state but also the mental states of agents. Unlike classical planning, epistemic planning explicitly characterizes and tracks how knowledge and beliefs evolve through actions. This capability enables planning in nondeterministic, partially observable, and multi-agent domains, where agents may reason about others' knowledge and beliefs. MEP plays important roles in applications such as games (De Giacomo et al. 2016), human-agent interaction (Baral et al. 2017), and cognitive science (Pietarinen 2003). Most existing approaches for MEP are based on Dynamic Epistemic Logic (DEL). In the DEL approach, both the epistemic states and actions (events) are represented using pointed Kripke models, and state transitions with respect to actions are computed through the *product update*, which integrates the epistemic model with the event model (Bolander and Andersen 2011). The DEL approach offers several

advantages for modeling MEP: it provides high expressivity for representing epistemic concepts such as nested beliefs and common knowledge, and it allows capturing complex epistemic dynamics arising from diverse action types and effects. In addition, the product update offers a uniform mathematical method for computing how agents' knowledge evolves in response to actions.

However, despite its expressiveness, the DEL approach faces several limitations. As a semantic approach, it requires a complete and explicit construction of both the initial model and the event model, which often demands expertise in modal logic. In complex domains, the automatic generation of these models is challenging. Moreover, it has been shown that MEP in DEL is undecidable, mainly due to the potential for unbounded growth of nested knowledge (Bolander and Andersen 2011). While decidable fragments exist, such as planning with propositional actions or pure epistemic actions (Löwe, Pacuit, and Witzel 2011; Yu, Wen, and Liu 2013; Bolander, Jensen, and Schwarzenruber 2015), strong restrictions are imposed that limit the expressivity.

To address the drawbacks of the semantic-based DEL approach, alternative syntactic approaches have also been explored. (Aucher 2011, 2012) investigates epistemic planning using DEL-sequent calculus and projection. Unlike the DEL approach, these works do not explicitly consider action sequences but instead produce a formal specification that the transition must fulfill. *Situation Calculus* (McCarthy and Hayes 1981; Reiter 2001), which is a first-order logical framework originally designed for physical actions, is used for relevant tasks such as belief change or progression (Fang and Liu 2013; Fang, Liu, and Wen 2015). However, to express knowledge and beliefs in the original design of the situation calculus requires second-order quantification, which makes it difficult to design a general planner but still preserve decidability. Via a modal variant of the epistemic situation calculus, (Belle and Lakemeyer 2014) presents a regression framework for multi-agent scenarios, which reduces queries after actions to the initial state so that tools for theorem proving can be applied. However, their approach does not handle MEP directly, and also suffers from limited expressivity, as it can only represent actions with objective pre- and postconditions, and cannot express announcements or nondeterministic actions.

To further extend the prior works and overcome the re-

striction on action theories, we investigate the cognitive-science concept *Theory of Mind* (ToM) - the cognitive mechanism by which one predicts other people’s actions by inferring their perceptions, beliefs, and desires (Ho, Saxe, and Cushman 2022). Higher-order ToM typically corresponds to more sophisticated cognitive and reasoning abilities, involving recursive mental-state attribution (Peloquin 2023). Inspired by the hierarchical structure of ToM, we propose a formalism of collective action theories, where each agent has an individual action theory that amounts to ToM of a certain order, and beliefs about other agents’ theories, which amount to lower-order ToMs. By preserving *bounded-order ToM* (BoToM) throughout the regression, our framework ensures the decidability of MEP represented within this formalism.

The rest of the paper is organized as follows. We begin by formalizing the logic for multi-agent epistemic reasoning in dynamic systems. Next, Section 3 introduces our hierarchical formalism for collective theories. We then discuss the various action types that can be represented within our framework. Subsequently, we present the regression approach in Section 5. Section 6 provides examples to demonstrate the expressiveness of our framework. In Section 7, we define MEP problems via logical entailment and prove their decidability. We propose an algorithm which solves MEPs and returns optimal solutions.

2 The Logic \mathcal{MS}

We consider a propositional variant of the logic \mathcal{ES}_n (Belle and Lakemeyer 2014): Given finite sets of agents Ag , actions Ac and atomic propositions Ap , the logic $\mathcal{MS}(\text{Ag}, \text{Ac}, \text{Ap})$ includes the following formulae:¹

$$\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{K}_i\varphi \mid [a]\varphi \mid \Box\varphi, \quad (1)$$

where $p \in \text{Ap}$. Other connectives like $\vee, \rightarrow, \equiv$ are used as the usual abbreviations. $\mathbf{K}_i\varphi$ means “agent i knows/believes φ ”,² $[a]\varphi$ is read as “ φ holds after action a ” and $\Box\varphi$ means “ φ holds after any sequence of actions”. For $G \subseteq \text{Ag}$, we use $\mathbf{E}_G\varphi$ as the abbreviation of $\bigwedge_{i \in G} \mathbf{K}_i\varphi$ and $\mathbf{E}_G^k\varphi$ means $\mathbf{E}_G(\dots(\mathbf{E}_G\varphi))$ with k nested \mathbf{E}_G . When $G = \text{Ag}$, we write $\mathbf{E}\varphi$ and $\mathbf{E}^k\varphi$. Note that \mathcal{ES}_n is a first-order logic with a countably infinite set of objects and actions. We restrict our language to a propositional fragment for decidable MEP. A first-order extension is certainly possible, but we leave it for future work. In addition, we include two “predicates”:

- $\text{Poss}(a)$: a is executable (in an agents’ mind);
- $I(a, a')$: a' is indistinguishable from a (for an agent).

Essentially, these predicates are interpreted as propositions in Ap , e.g., $\text{Poss}(\text{move})$ as poss.move . Let \mathcal{Z} denote the set of finite action sequences, i.e. $\mathcal{Z} = \text{Ac}^*$. We use $\langle \rangle$ to denote the empty one. A formula is *bounded* if it does not mention \Box , and it is *static* if it does not mention \Box or any $[a]$. We say a formula is *objective* if it mentions no \mathbf{K} operators and a formula is *subjective* if every proposition occurs

¹When the context is clear, we simply write \mathcal{MS} .

²“Knowledge” and “belief” are used interchangeably.

in the scope of a \mathbf{K} operator. A formula is *i-objective* if every occurrence of \mathbf{K}_i is in the scope of a \mathbf{K}_j where $j \neq i$, and a formula is *i-subjective* if every occurrence of propositions or \mathbf{K}_j for $j \neq i$ is in the scope of a \mathbf{K}_i . In addition, we use \top and \perp to represent tautologies and falsities. For $z = a_1 \cdots a_m$, we write $[z]\varphi$ instead of $[a_1] \cdots [a_m]\varphi$. Executability of sequence $\text{Exec}(z)$ is recursively defined as

- $\text{Exec}(\langle \rangle) := \top$;
- For $z = a \cdot \sigma$, $\text{Exec}(z) := \text{Poss}(a) \wedge [a]\text{Exec}(\sigma)$.

To evaluate the depth of a formula, we follow (Feng et al. 2025) and define:

Definition 1 (*i*-depth). For $i \in \text{Ag}$, the *i*-depth of a formula α , written as $\text{dep}[\alpha, i]$, is recursively defined as

- $\text{dep}[p, i] = 0$ for $p \in \text{Ap}$;
- $\text{dep}[\neg\alpha, i] = \text{dep}[\alpha, i]$;
- $\text{dep}[\alpha \wedge \beta, i] = \max(\text{dep}[\alpha, i], \text{dep}[\beta, i])$;
- $\text{dep}[[a]\alpha, i] = \text{dep}[\alpha, i]$ for $a \in \text{Ac}$;
- $\text{dep}[\Box\alpha, i] = \text{dep}[\alpha, i]$;
- $\text{dep}[\mathbf{K}_i\alpha, i] = \max(\max\{\text{dep}[\alpha, j] \mid j \neq i\} + 1, \text{dep}[\alpha, i])$;
- $\text{dep}[\mathbf{K}_j\alpha, i] = 0$ for $j \neq i$.

Intuitively, $\text{dep}[\alpha, i]$ measures the degree of nesting of agent i ’s mental state. For example, a formula $p \wedge \mathbf{K}_A \neg \mathbf{K}_B p$ has A -depth 2 and B -depth 0, as it involves what A believes that B believes, but not the actual beliefs of B . The depth of a formula α is defined as the maximal i -depth, i.e. $\text{dep}[\alpha] = \max_{i \in \text{Ag}} \text{dep}[\alpha, i]$.

The semantics of the logic involves several terms: *world*, *structure*, and *epistemic state*. Similar notions are in (Belle and Lakemeyer 2014, 2015; Feng et al. 2025). A world is defined as a mapping $w: \mathcal{Z} \rightarrow 2^{\text{Ap}}$. We use \mathcal{W} to denote the set of all worlds. Let $|\text{Ag}| = n$, for $k \geq 0$, we inductively define the epistemic states and structures as follows:

- $\mathcal{E}^0 = \{\{\emptyset\}\}$
- $\mathcal{S}^{k+1} = \{(w, e_1, \dots, e_{n-1}) \mid w \in \mathcal{W}, e_j \in \mathcal{E}^k \text{ for all } j\}$
- $\mathcal{E}^{k+1} = \wp(\mathcal{S}^{k+1})$, i.e. \mathcal{E}^{k+1} is the power set of \mathcal{S}^{k+1}

We call $s^k \in \mathcal{S}^k$ a k -structure and $e^k \in \mathcal{E}^k$ a k -state. We call k the depth of s^k and e^k . Intuitively, a structure consists of what an agent considers possible about the world and about other agents’ beliefs. When the context is clear, we omit the superscript k . By a model, we mean a tuple (w, e_1, \dots, e_n) (also written as (w, \vec{e})), where w is a world and e_i is the state of agent i . Note the difference between structures and models: a structure consists of $n - 1$ states of the same depth, while a model consists of n states whose depths can be different. For state \hat{e} and structure $s = (w, e_1, \dots, e_{n-1})$, we use $s \cup_i \hat{e}$ to denote the model obtained by inserting \hat{e} at the i -th position, i.e. $(w, e_1, \dots, e_{i-1}, \hat{e}, e_i, \dots, e_{n-1})$.

The satisfaction relation is defined as follows:

- $(w, \vec{e}) \models p$ iff $p \in w[\langle \rangle]$
- $(w, \vec{e}) \models \neg\alpha$ iff it is not the case that $(w, \vec{e}) \models \alpha$
- $(w, \vec{e}) \models \alpha \wedge \beta$ iff $(w, \vec{e}) \models \alpha$ and $(w, \vec{e}) \models \beta$
- $(w, \vec{e}) \models \mathbf{K}_i\alpha$ iff $s \cup_i e_i \models \alpha$ for all $s \in e_i$

We introduce the following notions. Essentially, they extend the semantics of $Poss$ and I to sequences:

Definition 2. We define $\Pi_w \subseteq \mathcal{Z}$ and $\sim_w \subseteq \mathcal{Z} \times \mathcal{Z}$ as:

- $\langle \rangle \in \Pi_w$. For $z = \sigma \cdot a$, $z \in \Pi_w$ iff $\sigma \in \Pi_w$ and $Poss(a) \in w[\sigma]$.
- $\langle \rangle \sim_w z$ iff $z = \langle \rangle$. For $z = \sigma \cdot a$ and $z' = \sigma' \cdot a'$, $z \sim_w z'$ iff $\sigma \sim_w \sigma'$ and $I(a, a') \in w[\sigma']$.

To characterize the situation after actions, we extend the world/state progression (Claßen 2013; Liu and Feng 2023) to multi-agent cases:

Definition 3. For $w \in \mathcal{W}$ and $z \in \mathcal{Z}$, the progression of w wrt z , written as w_z , is defined as the world w' such that for all $z' \in \mathcal{Z}$, $w[z \cdot z'] = w'[z']$. For epistemic state $e \in \mathcal{E}^k$, the progression of e wrt z , written as e_z , is defined as follows:

- When $k = 1$, i.e. e is a state of depth 1,
 $e_z = \{(w_\sigma, \{\emptyset\}), \dots, \{\emptyset\}) \mid (w, \{\emptyset\}), \dots, \{\emptyset\}) \in e,$
 $\sigma \in \Pi_w, z \sim_w \sigma\}$
 - When $k > 1$, i.e. e is a state deeper than 1,
 $e_z = \{(w_\sigma, e'_{1,\sigma}, \dots, e'_{n-1,\sigma}) \mid (w, e'_1, \dots, e'_{n-1}) \in e,$
 $\sigma \in \Pi_w \text{ and } z \sim_w \sigma\}$
- Where $e'_{j,\sigma}$ is the progression of e'_j wrt σ .

Now we can consider formulae with action modalities:

- $(w, \vec{e}) \models [a]\varphi$ iff $(w_a, \vec{e}_a) \models \varphi$
- $(w, \vec{e}) \models \Box\alpha$ iff $(w_z, \vec{e}_z) \models \alpha$ for all $z \in \mathcal{Z}$

We say a model and a formula α is *compatible* if, for all $i \in \text{Ag}$, the i -depth of α is not deeper than the epistemic state of agent i . Given a finite set $\Sigma \subseteq \mathcal{MS}$,³ $\alpha \in \mathcal{MS}$, we say Σ entails α (written as $\Sigma \models \alpha$) iff for every model (w, \vec{e}) compatible with each formula in $\Sigma \cup \{\alpha\}$, if $(w, \vec{e}) \models \varphi$ for all $\varphi \in \Sigma$, then $(w, \vec{e}) \models \alpha$. We say α is *valid* (written as $\models \alpha$) iff $\{\} \models \alpha$. We say a formula α is *satisfiable* if there is a compatible model (w, \vec{e}) s.t. $(w, \vec{e}) \models \alpha$. When α is objective, we write $w \models \alpha$ instead of $(w, \vec{e}) \models \alpha$. When α is i -subjective, we write $e_i \models \alpha$. Note that \mathcal{MS} cannot handle formulae of unbounded depth. When Σ is an infinite set of formulae with unbounded depth, i.e. for all k it exists $\psi \in \Sigma$ s.t. $dep[\psi] > k$, then $\Sigma \models \alpha$ is not well-defined.

The logic satisfies the $K45_n$ properties after any actions.

Theorem 1. For any $\alpha, \beta \in \mathcal{MS}$ and $i \in \text{Ag}$:

- (Prop) $\models \Box K_i \alpha$ if α is a (propositional) tautology;
- (Dist) $\models \Box K_i \alpha \wedge K_i (\alpha \rightarrow \beta) \rightarrow K_i \beta$;⁴
- (4) $\models \Box K_i \alpha \rightarrow K_i K_i \alpha$;
- (5) $\models \Box \neg K_i \alpha \rightarrow K_i \neg K_i \alpha$.

The static entailment amounts to $K45_n$ theorem proving.

Theorem 2. Given static Σ and α , $\Sigma \models \alpha$ iff α is derivable from Σ under the $K45_n$ axiom system.

³ Σ can also be an infinite set, provided that the depth of formulae is bounded, i.e. it exists $k \in \mathbb{N}$ s.t. $dep[\varphi] \leq k$ for all $\varphi \in \Sigma$.

⁴Throughout the paper, we assume that the modality \Box always has the lowest precedence among all operators.

It is proven by constructing a bisimulation between \mathcal{MS} models and $K45_n$ Kripke models. We refer to (Feng et al. 2025) for the proof.

With the syntax and semantics of the logic \mathcal{MS} , we now turn to multi-agent action theories, which specify the transition of not only the world but also agents' mental states.

3 Multi-Agent Action Theories

The proposed action theories are inspired by the BHL framework (Bacchus, Halpern, and Levesque 1999). As in BHL, a theory includes *precondition axioms*, *successor-state axioms* (SSAs), which describe the effect of actions, and *observation-indistinguishability* (OI) axioms, which distinguish the action actually being executed from those likely being executed. In addition, we include axioms on other agents' knowledge of actions.

Definition 4 (Individual Action Theory). For $i \in \text{Ag}$ and $k \in \mathbb{N}$, We define the k -order individual action theory (IAT) of agent i , denoted as $\Sigma^{i,k}$, as the following axioms:

- $\Sigma_{pre}^{i,k}$. For each action $a \in \text{Ac}$,
 $\Box Poss(a) \equiv \pi_a$ (2)
- $\Sigma_{ssa}^{i,k}$. For each $a \in \text{Ac}$ and $p \in \text{Ap}$,⁵
 $\Box [a]p \equiv \gamma_{p,a}$ (3)
- $\Sigma_{oi}^{i,k}$. For each $a, a' \in \text{Ac}$,
 $\Box I(a, a') \equiv \phi_{a,a'}$ (4)

- $\Sigma_{iat}^{i,k}$. For each $j \neq i$, j 's knowledge of actions $K_j \Sigma^{j,k-1}$ is included. When $k = 0$, $\Sigma^{j,k-1} := \top$. When $k > 0$, $\Sigma^{j,k-1}$ is an IAT of order $k - 1$,

where π_a , $\gamma_{p,a}$ and $\phi_{a,a'}$ are static formulae with a depth at most k . Note that for $k > 1$, $\Sigma^{j,k-1}$ includes other agents' knowledge about actions. The nesting continues until $k = 0$. Hence, a hierarchical structure is formed. Based on the IAT of each agent, we define the *collective action theory* (CAT):

Definition 5 (Collective Action Theory). For $k \in \mathbb{N}$, a k -order CAT, written as Γ , is defined as

$$\Gamma := \Gamma_{pre} \cup \Gamma_{ssa} \cup \Gamma_{oi} \cup \bigcup_{i \in \text{Ag}} K_i \Sigma^{i,k_i} \quad (5)$$

where Γ_{pre} , Γ_{ssa} and Γ_{oi} are respectively of the form (2-4) and the corresponding π_a , $\gamma_{p,a}$ and $\phi_{a,a'}$ are static formulae of depth at most k . Σ^{i,k_i} is the IAT of agent i . We use $\text{Ord}(\Gamma, i) := k_i$ to denote the order of agent i 's action theory. $\text{Ord}(\Gamma)$ denotes the maximal order of Γ . By a Γ -model, we mean a model (w, \vec{e}) compatible with Γ , and $(w, \vec{e}) \models \Gamma$.

Essentially, a k -order IAT corresponds to the mind of an agent with k -order ToM. There are two reasons for having such a hierarchy: The first is to distinguish an agent's IAT from what other agents believe to be his or her IAT, and the second is for the hierarchical nature of ToM: A child with 1-order ToM could accurately reason about others' beliefs about the world state, and can distinguish what is believed

⁵Propositions about $Poss$ or I , e.g. $poss_move$, are excluded.

by others from what is true. When two children with 1-order ToM share their rationale behind actions, can they accurately model the knowledge about the mental state transition of the other? The answer is no. A k -order agent cannot reason about the knowledge of another k -order agent, as representing others' perspectives consumes one level of nesting.

Definition 6. Given theory Γ which includes $\mathbf{K}_i \Sigma^{i,k}$ and $\Sigma^{i,k}$ of the form $\Sigma_{pre}^{i,k} \cup \Sigma_{ssa}^{i,k} \cup \Sigma_{oi}^{i,k} \cup \bigcup_{j \neq i} \mathbf{K}_j \Sigma_j^{j,k-1}$, the *slice of Γ wrt agent i* , written as Γ^i , is a CAT defined as

$$\Gamma^i := \Sigma_{pre}^{i,k} \cup \Sigma_{ssa}^{i,k} \cup \Sigma_{oi}^{i,k} \cup \mathbf{K}_i \Sigma^{i,k} \cup \bigcup_{j \neq i} \mathbf{K}_j \Sigma_j^{j,k-1}$$

i.e. the precondition axioms, SSAs and OI axioms of Γ^i are the corresponding parts in $\Sigma^{i,k}$.

Γ^i is a CAT reconstructed via agent i 's IAT in Γ . It is not difficult to prove $(\Gamma^i)^i = \Gamma^i$.

Now we consider the existence of Γ -models: For any model (w, \vec{e}) , since Γ only regulates the state transition but not the initial state, one can construct a Γ -model wrt (w, \vec{e}) which satisfies the same static formulae:

Theorem 3. Given theory Γ and model (w, \vec{e}) compatible with Γ , there exists a Γ -model $(w_\Gamma, e_{1,\Gamma}, \dots, e_{n,\Gamma})$ (also written as $(w_\Gamma, \vec{e}_\Gamma)$) of the same depth such that for each compatible static formula φ :

$$(w, \vec{e}) \models \varphi \text{ iff } (w_\Gamma, \vec{e}_\Gamma) \models \varphi \quad (6)$$

Proof. Given (w, \vec{e}) , we recursively construct $(w_\Gamma, \vec{e}_\Gamma)$ such that for all $p \in \text{Ap}$, $z \in \mathcal{Z}$ and $a, a' \in \text{Ac}$,

- $w_\Gamma[\langle \rangle] = w[\langle \rangle]$;
- $p \in w_\Gamma[z \cdot a]$ iff $(w_\Gamma, \vec{e}_\Gamma) \models [z] \gamma_{p,a}^\Gamma$;
- $\text{Poss}(a) \in w_\Gamma[z]$ iff $(w_\Gamma, \vec{e}_\Gamma) \models [z] \pi_a^\Gamma$;
- $I(a, a') \in w_\Gamma[z]$ iff $(w_\Gamma, \vec{e}_\Gamma) \models [z] \phi_{a,a'}^\Gamma$;
- For any $i \in \text{Ag}$, $e_{i,\Gamma} = \{s_{\Gamma^i} \mid s \in e_i\}$, where
For $s = (w, e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_n)$,
 $s_{\Gamma^i} = (w_{\Gamma^i}, e_{1,\Gamma^i}, \dots, e_{i-1,\Gamma^i}, e_{i+1,\Gamma^i}, \dots, e_{n,\Gamma^i})$.

Here $\gamma_{p,a}^\Gamma$, π_a^Γ and $\phi_{a,a'}^\Gamma$ are respectively the right-hand side of Γ_{pre} , Γ_{ssa} and Γ_{oi} . For all e_j and w , we construct e_{j,Γ^i} and w_{Γ^i} in a similar way but wrt Γ^i instead of Γ . We prove Eq. 6 via a double-induction: The outer induction is on the order of Γ and the inner induction is on the construction of formula φ .

Base case 1: For a zero-order CAT, it reduces to the single-agent cases with objective φ . We refer to (Lakemeyer and Lesvesque 2011) for the proof details.

I.H.1: Suppose that Eq.(6) holds for all Γ' which satisfies $\text{Ord}(\Gamma') < k$, formula φ , and (w, \vec{e}) compatible with Γ' .

Base case 2: For any Γ with $\text{Ord}(\Gamma) = k$, It is trivial that $(w, \vec{e}) \models p$ iff $(w_\Gamma, \vec{e}_\Gamma) \models p$.

I.H.2: Suppose that $(w, \vec{e}) \models \alpha$ iff $(w_{\Gamma'}, \vec{e}_{\Gamma'}) \models \alpha$ for all Γ' with $\text{Ord}(\Gamma') = k$ and (w, \vec{e}) compatible with Γ' . Now we consider a fixed Γ . Induction on Boolean operators is trivial.

For induction on the modality \mathbf{K}_i ,

$$\begin{aligned} (w, \vec{e}) \models \mathbf{K}_i \alpha & \\ \Leftrightarrow \text{f.a. } s \in e_i, s \cup_i e_i \models \alpha & \quad (\text{Semantics}) \\ \Leftrightarrow \text{f.a. } s \in e_i, s_{\Gamma^i} \cup_i e_{i,\Gamma^i} \models \alpha & \quad (\#) \\ \Leftrightarrow \text{f.a. } s' \in e_{i,\Gamma}, s' \cup_i e_{i,\Gamma} \models \alpha & \quad (e_{i,\Gamma} = e_{i,\Gamma^i}) \\ \Leftrightarrow e_{i,\Gamma} \models \mathbf{K}_i \alpha \Leftrightarrow (w_\Gamma, \vec{e}_\Gamma) \models \mathbf{K}_i \alpha & \end{aligned}$$

There are two cases for (#): When $\text{Ord}(\Gamma^i) < k$, then I.H.1 is applied. Otherwise, we apply I.H. 2 as $\text{Ord}(\Gamma^i) = k$.

To prove that $(w_\Gamma, \vec{e}_\Gamma)$ is indeed a Γ -model, we check each entry of Γ . For the precondition, SSAs and OI axioms, the proof applies simply the semantics and the definition of $(w_\Gamma, \vec{e}_\Gamma)$. For the IAT of each agent i :

$$\begin{aligned} (w_\Gamma, \vec{e}_\Gamma) \models \mathbf{K}_i \Sigma^{i,k_i} & \\ \Leftrightarrow s' \cup_i e_{i,\Gamma} \models \Sigma^{i,k_i} \text{ f.a. } s' \in e_{i,\Gamma} & \\ \Leftrightarrow s' \cup_i e_{i,\Gamma} \models \Sigma^{i,k_i} \wedge \mathbf{K}_i \Sigma^{i,k_i} \text{ f.a. } s' \in e_{i,\Gamma} & \\ \Leftrightarrow s' \cup_i e_{i,\Gamma} \models \Gamma^i \text{ f.a. } s' \in e_{i,\Gamma} & \\ \Leftrightarrow s_{\Gamma^i} \cup_i e_{i,\Gamma^i} \models \Gamma^i \text{ f.a. } s \in e_i & \end{aligned}$$

It is equivalent to prove that for each $s \in e_i$, $s_{\Gamma^i} \cup_i e_{i,\Gamma^i}$ is a Γ^i -model, which can be proved via induction. \square

We define a knowledge base with bounded-order ToM (BoToM-KB) as follows:

Definition 7 (BoToM-KB). A finite set of formulae $\Sigma_0 \cup \Gamma$ forms a BoToM-KB, if there is $k \in \mathbb{N}$ such that Σ_0 is a static formula of depth at most k , and Γ is a k -order CAT.

Σ_0 can be seen as the initial state in the DEL approach, and to some extent, Γ corresponds to the DEL state-transition system. However, to represent the ToM hierarchy in DEL, the event models and product update need to be modified such that the depths of the Kripke states are considered. So far, no one has investigated such modification.

4 Action Types

Three types of action are usually considered in multi-agent domains: *ontic actions*, *sensing*, and *announcements*. We show how these actions can be modeled in our framework.

Ontic actions have actual effects on the world state. Given a theory Γ , an action a is ontic if there is a $p \in \text{Ap}$ such that Γ does not entail $\Box \gamma_{p,a} \equiv p$. Our framework can represent ontic action with uncertain or unexpected effects. Uncertain effects, either caused by the partial observability of the agent or the non-deterministic nature of the world, can be modeled in the sense that $\phi_{a,a'}$ of a fixed a holds for more than one a' . Unexpected effects are modeled in the sense that the action which actually happens is considered impossible, i.e. $\phi_{a,a}$ does not hold.

In contrast to ontic actions, sensing and announcements do not change the actual world but only what agents believe. Thus, these actions only have trivial SSAs with $\gamma_{p,a} \equiv p$. Let see_h be sensing the tossing result Head, then a precise sensing can be represented by

$$\begin{aligned} \Box \text{Poss}(see_h) &\equiv Hd; & \Box I(see_h, see_h) &\equiv \top; \\ \Box I(see_h, a') &\equiv \perp \text{ for } a' \neq see_h \end{aligned}$$

Let Σ^i include these axioms, then $\mathbf{K}_i \Sigma^i \models [\text{see_h}] \mathbf{K}_i Hd$. Our framework supports announcement to any group of agents. Let $G!Hd$ denote the announcement to $G \subseteq \text{Ag}$ that the coin shows head. For each $i \in G$, let $\Sigma^{i,k}$ include

$$\begin{aligned} \square Poss(G!Hd) &\equiv Hd; & \square I(G!Hd, G!Hd) &\equiv \top; \\ \square I(G!Hd, a') &\equiv \perp \text{ for } a' \neq G!Hd. \end{aligned}$$

Let Γ include $\mathbf{K}_i \Sigma^{i,k}$ for each $i \in G$, then

$$\Gamma \models [G!Hd] E_G^l Hd$$

for any l up to the order of theory Γ . Other forms of actions, such as noisy sensing or deceptive announcements, will be demonstrated by examples in a later section.

In fact, the notion of ‘‘actions’’ here is similar to ‘‘events’’ in the DEL approach, as the executor of an action is not specified. While complex constructs such as action profiles (multiple agents simultaneously choose actions) can naturally be expressed within our framework, we focus on the general formalism and do not elaborate on these constructs.

5 Regression

For the action theories mentioned above, we introduce the regression method, which transforms epistemic queries on a future situation to those on the current situation.

Definition 8 (Regressable). *We call a formula φ regressable wrt Γ iff φ is bounded, and $\text{dep}[\varphi, i] \leq \text{Ord}(\Gamma, i)$ for all i .*

Definition 9 (Regression function). *For $z \in \mathcal{Z}$ and k -order collective theory Γ , we recursively define $\mathcal{R}[\Gamma, z, \alpha]$:*

- $\mathcal{R}[\Gamma, \langle \rangle, p] := p$ for p not mentioning $Poss$ or I ;
- $\mathcal{R}[\Gamma, \sigma \cdot a, p] := \mathcal{R}[\Gamma, \sigma, \gamma_{p,a}^\Gamma]$;
- $\mathcal{R}[\Gamma, \sigma, Poss(a)] := \mathcal{R}[\Gamma, \sigma, \pi_a^\Gamma]$;
- $\mathcal{R}[\Gamma, \sigma, I(a, a')] := \mathcal{R}[\Gamma, \sigma, \phi_{a,a'}^\Gamma]$;
- $\mathcal{R}[\Gamma, z, \alpha \wedge \beta] := \mathcal{R}[\Gamma, z, \alpha] \wedge \mathcal{R}[\Gamma, z, \beta]$;
- $\mathcal{R}[\Gamma, z, \neg \alpha] := \neg \mathcal{R}[\Gamma, z, \alpha]$;
- $\mathcal{R}[\Gamma, z, [a]\alpha] := \mathcal{R}[\Gamma, z \cdot a, \alpha]$;
- $\mathcal{R}[\Gamma, \langle \rangle, \mathbf{K}_i \varphi] := \mathbf{K}_i (\mathcal{R}[\Gamma^i, \langle \rangle, \varphi])$;
- $\mathcal{R}[\Gamma, \sigma \cdot a, \mathbf{K}_i \varphi] := \mathcal{R}[\Gamma, \sigma, \mathbf{K}_i \bigwedge_{a'} ((Poss(a') \wedge I(a, a')) \rightarrow [a'] \varphi)]$

The regression always terminates and returns a static formula of size $O(|\varphi| \cdot |\Gamma| \cdot |\text{Ac}|^{k \cdot m})$, where k is the order of Γ and m is the action length of $[z]\varphi$. When the second argument is an empty sequence, we write $\mathcal{R}[\Gamma, \alpha]$ instead of $\mathcal{R}[\Gamma, \langle \rangle, \alpha]$.

The main motivation for introducing hierarchical structures and descending orders of action theories is to limit the depth of regression results. Via induction, we prove that

Theorem 4. *For any theory Γ , φ regressable wrt Γ , $\text{dep}[\mathcal{R}[\Gamma, z, \varphi], i] \leq \text{dep}[\Gamma, i]$ for all z and $i \in \text{Ag}$.*

Namely, the regression never results in a formula deeper than the order of Γ . The regression is correct in the sense that any Γ -model satisfies α after sequence z if and only if the model satisfies the initial condition $\mathcal{R}[\Gamma, z, \alpha]$:

Lemma 1. *Given action theory Γ , for all Γ -model (w, \vec{e}) , $z \in \mathcal{Z}$ and static formula α regressable wrt Γ ,*

$$(w, \vec{e}) \models [z]\alpha \text{ iff } (w, \vec{e}) \models \mathcal{R}[\Gamma, z, \alpha] \quad (7)$$

Proof sketch: We proceed with a triple-induction: The outer-induction is on the order of the theory Γ , the mid-induction is on the length of the action sequence z , and the inner-induction is on the construction of formula α .

The lemma can be easily extended to regression on arbitrary bounded formulae:

Theorem 5. *Given action theory Γ , for all Γ -model (w, \vec{e}) , $z \in \mathcal{Z}$ and bounded formula α regressable wrt Γ ,*

$$(w, \vec{e}) \models [z]\alpha \text{ iff } (w, \vec{e}) \models \mathcal{R}[\Gamma, z, \alpha] \quad (8)$$

Now we are ready for the regression theorem:

Theorem 6 (Regression). *Given BoToM-KB $\Sigma_0 \cup \Gamma$, for all $z \in \mathcal{Z}$, and formula α regressable wrt Γ ,*

$$\Sigma_0 \cup \Gamma \models [z]\alpha \text{ iff } \Sigma_0 \models \mathcal{R}[\Gamma, z, \alpha] \quad (9)$$

Proof. It suffices to prove the equivalence of the statements:

$$\text{F.a. } (w, \vec{e}) \models \Gamma, \text{ if } (w, \vec{e}) \models \Sigma_0, \text{ then } (w, \vec{e}) \models [z]\alpha \quad (10)$$

$$\text{F.a. } (w, \vec{e}), \text{ if } (w, \vec{e}) \models \Sigma_0, \text{ then } (w, \vec{e}) \models \mathcal{R}[\Gamma, z, \alpha] \quad (11)$$

(10) \Rightarrow (11): For any model (w, \vec{e}) :

$$(w, \vec{e}) \models \Sigma_0 \Rightarrow (w_\Gamma, \vec{e}_\Gamma) \models \Sigma_0 \quad (\text{Thm. 3})$$

$$\Rightarrow (w_\Gamma, \vec{e}_\Gamma) \models [z]\alpha \quad (\text{Eq. 10})$$

$$\Rightarrow (w_\Gamma, \vec{e}_\Gamma) \models \mathcal{R}[\Gamma, z, \alpha] \quad (\text{Thm. 5})$$

$$\Rightarrow (w, \vec{e}) \models \mathcal{R}[\Gamma, z, \alpha] \quad (\text{Thm. 3})$$

(10) \Leftarrow (11): For any Γ -model $(w_\Gamma, \vec{e}_\Gamma)$,

$$(w_\Gamma, \vec{e}_\Gamma) \models \Sigma_0 \Rightarrow (w_\Gamma, \vec{e}_\Gamma) \models \mathcal{R}[\Gamma, z, \alpha] \quad (\text{Eq. 11})$$

$$\Rightarrow (w_\Gamma, \vec{e}_\Gamma) \models [z]\alpha \quad (\text{Thm. 5})$$

□

By Thm. 5 we proved the following result, which will later be used in the planning task.

Corollary 1. *For $z \in \mathcal{Z}$, and any φ, ψ regressable wrt Γ ,*

- *if $\models \varphi \equiv \psi$, then $\models \mathcal{R}[\Gamma, z, \varphi] \equiv \mathcal{R}[\Gamma, z, \psi]$*
- *if $\models \mathcal{R}[\Gamma, \varphi] \equiv \mathcal{R}[\Gamma, \psi]$, then $\models \mathcal{R}[\Gamma, z, \varphi] \equiv \mathcal{R}[\Gamma, z, \psi]$*

6 Examples

We formalize action theories for two examples: Sally-Anne Test (Wimmer and Perner 1983) and Die Rolling. Due to the space limit, we only present the key entries of the theories that are relevant to regression. In particular, we consider the dynamics of four actions:

- *move*: To move the marble. We consider the unexpected effect of the action, or as described in (Baral et al. 2022), an agent is oblivious of the actions’ occurrence.
- *roll*: To roll a die. A non-deterministic action.
- *show2*: An agent saw that the die showed two. We model noisy sensing, namely, a distracted agent might misread the number.
- *!bob_3*: Bob announces that the die shows three (even if it does not). An example of deceptive announcements.

Example 1 (Sally-Anne Test). *Initially, Sally’s marble is in a basket (and everyone knows that), then Anne (secretly) moves the marble into a box. In a test, the child will be asked ‘‘Will Sally look in the box for her marble?’’*

The test is a classic psychological experiment used to assess the order of ToM: If the participant believes that Sally will not look in the box, it indicates that the child has developed first-order ToM and can understand another person's beliefs, even if they are false. Otherwise, the child has at most zero-order ToM. We show how different orders of ToM can be characterized:

- $\text{Ag} = \{S, C_0, C_1\}$, which respectively stands for Sally, Child with 0-order ToM and with 1-order ToM.
- $\text{Ap} = \{InBt, InBx\}$
The marble is in the Basket / in the Box.
- $\text{Ac} = \{move, void, look_bx\}$, *void* for nothing happens, *look_bx* for “look in the box for the marble”

We consider a collective theory Γ which includes

- $\mathbf{K}_{C_0}\Sigma^{C_0}$, i.e. a 0-order IAT for child C_0 , and
- $\mathbf{K}_{C_1}\Sigma^{C_1}$, a 1-order IAT for child C_1 .

Σ^{C_0} includes

- $\Box Poss(look_bx) \equiv InBx$,
- $\Box[move]InBx \equiv \top$, $\Box[move]InBt \equiv \perp$, and
- $\Box I(a, a) \equiv \top$, $\Box I(a, a') \equiv \perp$ for all $a \neq a'$.

Σ^{C_1} differs from Σ^{C_0} by

- $\Box Poss(look_bx) \equiv \mathbf{K}_S InBx$, i.e. Sally will look in the box only when she believes that the marble is in the box.
- Σ^{C_1} includes $\mathbf{K}_S \Sigma^{S,0}$ where $\Sigma_{oi}^{S,0}$ includes $\Box I(move, void) \equiv \top$, $\Box I(move, move) \equiv \perp$. Namely, Sally is oblivious to the move.

Let Σ_0 be the initial state. We encode the children's answers to the test as their knowledge about the possibility of action *look_bx*. By regression, we prove that

$$\begin{aligned} \Sigma_0 \cup \Gamma &\models \mathbf{K}_{C_0}([move]Poss(look_bx)) \\ \Leftrightarrow \Sigma_0 &\models \mathcal{R}[\Gamma, \mathbf{K}_{C_0}([move]Poss(look_bx))] \\ \Leftrightarrow \Sigma_0 &\models \mathbf{K}_{C_0}(\mathcal{R}[\Gamma^{C_0}, move, Poss(look_bx)]) \\ \Leftrightarrow \Sigma_0 &\models \mathbf{K}_{C_0}(\mathcal{R}[\Gamma^{C_0}, move, \top]) \Leftrightarrow \Sigma_0 \models \top \end{aligned}$$

Namely, in any initial case, C_0 believes that Sally will look in the box for the marble. In contrast, child C_1 does not believe that Sally will look in the box. By regression,

$$\begin{aligned} \Sigma_0 \cup \Gamma &\models \neg \mathbf{K}_{C_1}([move]Poss(look_bx)) \\ \Leftrightarrow \Sigma_0 &\models \neg \mathbf{K}_{C_1}(\mathcal{R}[\Gamma^{C_1}, move, \mathbf{K}_S InBx]) \\ \Leftrightarrow \Sigma_0 &\models \neg \mathbf{K}_{C_1} \mathbf{K}_S(\mathcal{R}[\Gamma^S, void, InBx]) \\ \Leftrightarrow \Sigma_0 &\models \neg \mathbf{K}_{C_1} \mathbf{K}_S(InBx) \end{aligned}$$

According to the example, the child knows that Sally knows that initially, the marble is in the basket instead of the box.

Example 2 (Die Rolling). *Alice and Bob roll a die. The result can be shown to them, but when Alice is distracted, she may misread the number. In addition, Bob can announce the number, even if it is wrong, and Alice will believe it anyhow.*

- $\text{Ag} = \{A, B\}$, i.e. two agents Alice and Bob.

- $\text{Ac} = \{roll, roll2, show2, !bob_3, \dots\}$
roll is a nominal action and *roll1* to *roll6* are factual actions with the resulting number. *show2* and *!bob_3* are, respectively, the noisy sensing and announcement mentioned at the beginning of this section.
- $\text{Ap} = \{N1, \dots, N6, Dt\}$. *Dt* means Alice is distracted, *Nx* stands for the number that the die shows.

We consider the collective theory Γ which includes $\mathbf{K}_A \Sigma^A$ and $\mathbf{K}_B \Sigma^B$. Alice's individual theory Σ^A consists of

- $\Box Poss(show2) \equiv (Dt \rightarrow (N2 \vee N3)) \wedge (\neg Dt \rightarrow N2)$
Alice may misread 3 as 2 when she is distracted.
- $\Box Poss(!bob_3) \equiv N3 \wedge \neg N1 \wedge \dots \wedge \neg N6$
Alice always believes the truth of Bob's announcement.
- $\Box[roll2]N2 \equiv \top$ and $\Box[roll2]Nx \equiv \perp$ for $x \neq 2$.
- $\Box I(roll, rollx) \equiv \top$ for $x \in \{1, \dots, 6\}$

Bob's theory Σ^B differs from Alice's by

- $\Box Poss(show2) \equiv N2$, i.e. Bob never misreads.
- $\Box Poss(!bob_3) \equiv \top$, i.e. Bob can always announce that the number is 3, as long as he wants to do so.

After rolling the die and showing the result 2, to achieve a state where Alice does not know that the die shows 2, by regression we prove that the initial situation must fulfill

$$\Sigma_0 \cup \Gamma \models [roll][show2] \neg \mathbf{K}_A N2 \Leftrightarrow \Sigma_0 \models \neg \mathbf{K}_A \neg Dt$$

Namely, Alice cannot confirm that she stays focused. Lastly, we show how Bob can make a false announcement and manipulate Alice's beliefs:

$$\Sigma_0 \cup \Gamma \models [!bob_3](\mathbf{K}_A N3) \Leftrightarrow \Sigma_0 \models \top$$

Namely, Alice will always believe Bob's announcement.

7 Multi-agent Epistemic Planning

Multi-agent epistemic planning amounts to finding an action sequence such that, after taking the actions, a given goal is fulfilled. In the DEL approach, MEP is known to be robustly undecidable mainly because agents can reason with an arbitrary nesting of beliefs (Aucher and Bolander 2013). With the bounded-order hierarchy of the action theories, we restrict the degree of nesting and thus preserve decidability.

Definition 10 (BoToM-MEP). *Given a BoToM-KB $\Sigma_0 \cup \Gamma$ and a goal α regressable wrt Γ , the BoToM-MEP problem is to decide whether a finite sequence z exists such that*

$$\Sigma_0 \cup \Gamma \models Exec(z) \wedge [z]\alpha \quad (12)$$

A sequence z which fulfills (12) is called a *solution* of the MEP wrt Σ_0, Γ and α . Note that the DEL approach defines MEP based on the satisfaction of a single model and therefore requires complete specification of the initial state; and our definition is based on logical entailment, where multiple, even infinitely many models are involved. Thus, we accept incomplete specification.

Theorem 7. *For $\sigma_1, \sigma_2 \in \mathcal{Z}$ s.t. $\mathcal{R}[\Gamma, Exec(\sigma_1) \wedge [\sigma_1]\alpha]$ and $\mathcal{R}[\Gamma, Exec(\sigma_2) \wedge [\sigma_2]\alpha]$ are logically equivalent, if $z \cdot \sigma_1$ is a solution of the problem, then $z \cdot \sigma_2$ is also a solution.*

Proof. Suppose that $z \cdot \sigma_1$ is a solution,

$$\begin{aligned}
& \Sigma_0 \cup \Gamma \models Exec(z \cdot \sigma_1) \wedge [z \cdot \sigma_1]\alpha \\
& \iff \Sigma_0 \cup \Gamma \models Exec(z) \wedge [z]Exec(\sigma_1) \wedge [z \cdot \sigma_1]\alpha \\
& \iff \Sigma_0 \cup \Gamma \models \mathcal{R}[\Gamma, Exec(z)] \\
& \quad \wedge \mathcal{R}[\Gamma, z, Exec(\sigma_1) \wedge [\sigma_1]\alpha] \quad (\text{Thm. 6}) \\
& \iff \Sigma_0 \cup \Gamma \models \mathcal{R}[\Gamma, Exec(z)] \\
& \quad \wedge \mathcal{R}[\Gamma, z, Exec(\sigma_2) \wedge [\sigma_2]\alpha] \quad (\text{Cor. 1}) \\
& \iff \Sigma_0 \cup \Gamma \models \mathcal{R}[\Gamma, Exec(z) \wedge [z]Exec(\sigma_2) \wedge [z \cdot \sigma_2]\alpha] \\
& \iff \Sigma_0 \cup \Gamma \models Exec(z \cdot \sigma_2) \wedge [z \cdot \sigma_2]\alpha
\end{aligned}$$

□

Proposition 1. For $z = a \cdot \sigma$, $\Sigma_0 \models \mathcal{R}[\Gamma, Exec(z) \wedge [z]\alpha]$ if and only if $\Sigma_0 \models Poss(a) \wedge \mathcal{R}[\Gamma, a, \mathcal{R}[Exec(\sigma) \wedge [\sigma]\alpha]]$

With the above theorem and proposition, we are ready to prove the most important result of this work:

Theorem 8. *BoToM-MEP is decidable.*

Proof. Given a k -order theory Γ and α regressable wrt Γ . By Thm. 6, to decide Eq. 12 amounts to decide

$$\Sigma_0 \models \mathcal{R}[\Gamma, Exec(z) \wedge [z]\alpha] \quad (13)$$

Since both Σ_0 and the regression are static, it is reduced to $K45_n$ reasoning. By Thm. 4, the regression function always returns a formula with depth bounded by k . As pointed out in (Halpern 1995), for a fixed k and finite Ap , there are only finitely many inequivalent static formulae. Let the number of such formulae be \mathcal{N} . By Thm. 7 and the Pigeonhole Principle, if (12) has a solution, there must be a solution z of length at most \mathcal{N} . Since Ac is also finite, by enumerating all action sequences up to length \mathcal{N} , one can decide if a solution exists. □

In fact, it is not needed to enumerate all sequences or to compute the exact number of \mathcal{N} . It can be solved more efficiently by Algorithm 1: A queue of action sequence candidates Cand is maintained and initialized with the empty one. A dictionary State , which maps action sequences to \mathcal{MS} formulae, is maintained and initialized as $\{\langle \rangle : \alpha\}$. In each iteration of the while-loop, the algorithm pops out the first sequence z in Cand , generates a longer sequence $a \cdot z$ for each $a \in \text{Ac}$, and tests whether Σ_0 entails the regression.⁶ Line 13 tests if the regression is equivalent to a formula in State .⁷ If the regression wrt $a \cdot z$ is equivalent to one wrt an explored sequence z' , by Thm. 7 the shortest solution expanded from $a \cdot z$ will not be shorter as those expanded from z' . Hence, we can prune $a \cdot z$. Since the size of State will be at most \mathcal{N} , either a solution is found, or eventually every regression can be found in State , then every sequence in Cand will be pruned and it terminates when the queue is empty.

Theorem 9. *Algorithm 1 terminates and returns FALSE if the BoToM-MEP has no solution. When a solution exists, the algorithm returns a solution z such that for any other solution z' , $|z| \leq |z'|$.*

⁶By Prop.1 it suffices to perform a single step of regression.

⁷Though State could be huge, equivalence check can be efficient with proper knowledge compilation methods (Halpern 1995).

Algorithm 1: Breadth-First BoToM-MEP Search

```

1: function botommep( $\Sigma_0, \Gamma, \alpha$ ):
2:   Cand  $\leftarrow \{\langle \rangle\}$ 
3:   State $\{\langle \rangle\} \leftarrow \alpha$ 
4:   if  $\Sigma_0 \models \mathcal{R}[\Gamma, \alpha]$  then
5:     return  $\langle \rangle$ 
6:   end if
7:   while  $\text{Cand}$  not empty do
8:      $z \leftarrow \text{POP}(\text{Cand})$ 
9:      $\varphi \leftarrow \text{State}[z]$ 
10:    for  $a \in \text{Ac}$  do
11:      if  $\Sigma_0 \models \mathcal{R}[\Gamma, Poss(a) \wedge [a]\varphi]$  then
12:        return  $a \cdot z$ 
13:      else if  $\mathcal{R}[\Gamma, Exec(a) \wedge [a]\varphi] \notin \text{State}$  then
14:         $\text{State}[a \cdot z] \leftarrow \mathcal{R}[\Gamma, Poss(a) \wedge [a]\varphi]$ 
15:         $\text{PUSH}(\text{Cand}, a \cdot z)$ 
16:      end if
17:    end for
18:  end while
19: return FALSE
20: end function

```

8 Conclusion

Our work is the first attempt to solve MEP in the situation calculus. Compared with the de facto standard approach via DEL model update, our syntactic method avoids the manual construction of semantic models and enables incomplete specifications. With the price of not being able to deal with common knowledge or arbitrary depth of belief nesting,⁸ we preserve most of the expressiveness of the DEL approach and have decidable planning. To our knowledge, no existing framework can express all the action types we mentioned, while MEP remains decidable.

This framework brings many future directions: Strategy reasoning and epistemic protocol synthesis generate (group) strategies to ensure temporal goals (Xiong and Liu 2016; Aucher, Maubert, and Pinchinat 2014). In our framework, strategies and protocols can be characterized by additional preconditions. With a probabilistic extension, one can represent the likelihood of non-deterministic actions and measure an agent's degree of belief (Bacchus, Halpern, and Levesque 1999; Liu and Feng 2023; Feng and Lakemeyer 2024). It is also interesting to implement an actual epistemic planner (Wan, Fang, and Liu 2021; Liu and Liu 2018; Ware and Siler 2021a,b; Muise et al. 2022), or to build a connection with epistemic game theory and description languages for games (Thielscher 2017).

Acknowledgments

We thank Yongmei Liu (Sun Yat-sen University) for helpful discussions and suggestions regarding related work. This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) 2236/2 'UnRAVeL' and the EU ICT-48 2020 project TAILOR (No. 952215).

⁸We treat agents as bounded reasoners with finite-order ToMs.

References

- Aucher, G. 2011. DEL-sequents for progression. *Journal of Applied Non-Classical Logics*, 21(3-4): 289–321.
- Aucher, G. 2012. DEL-sequents for regression and epistemic planning. *Journal of Applied Non-Classical Logics*, 22(4): 337–367.
- Aucher, G.; and Bolander, T. 2013. *Undecidability in epistemic planning*. Ph.D. thesis, INRIA.
- Aucher, G.; Maubert, B.; and Pinchinat, S. 2014. Automata techniques for epistemic protocol synthesis. *arXiv preprint arXiv:1404.0844*.
- Bacchus, F.; Halpern, J. Y.; and Levesque, H. J. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *Artificial Intelligence*, 111(1-2): 171–208.
- Baral, C.; Bolander, T.; van Ditmarsch, H.; and McIlraith, S. 2017. Epistemic planning (Dagstuhl seminar 17231). *Dagstuhl Reports*, 7(6): 1–47.
- Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2022. An action language for multi-agent domains. *Artificial Intelligence*, 302: 103601.
- Belle, V.; and Lakemeyer, G. 2014. Multiagent only knowing in dynamic systems. *Journal of Artificial Intelligence Research*, 49: 363–402.
- Belle, V.; and Lakemeyer, G. 2015. Semantical considerations on multiagent only knowing. *Artificial Intelligence*, 223: 1–26.
- Bolander, T.; and Andersen, M. B. 2011. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1): 9–34.
- Bolander, T.; Jensen, M. H.; and Schwarzentruher, F. 2015. Complexity results in epistemic planning. In *24th International Joint Conference on Artificial Intelligence*, 2791–2797. AAAI Press.
- Claßen, J. 2013. *Planning and verification in the agent language Golog*. Ph.D. thesis, Aachen, Techn. Hochsch., Diss., 2013.
- De Giacomo, G.; Murano, A.; Rubin, S.; Di Stasio, A.; et al. 2016. Imperfect-Information Games and Generalized Planning. In *IJCAI*, 1037–1043.
- Fang, L.; and Liu, Y. 2013. Multiagent knowledge and belief change in the situation calculus. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 304–312.
- Fang, L.; Liu, Y.; and Wen, X. 2015. On the Progression of Knowledge and Belief for Nondeterministic Actions in the Situation Calculus. In *IJCAI*, 2955–2963.
- Feng, Q.; and Lakemeyer, G. 2024. Probabilistic Multi-agent Only-Believing. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 571–579.
- Feng, Q.; Wilk, H.; Khan, S. M.; and Lakemeyer, G. 2025. Translating Multi-Agent Modal Logics of Knowledge and Belief into Decidable First-Order Fragments. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems*, 740–748.
- Halpern, J. Y. 1995. The effect of bounding the number of primitive propositions and the depth of nesting on the complexity of modal logic. *Artificial Intelligence*, 75(2): 361–372.
- Ho, M. K.; Saxe, R.; and Cushman, F. 2022. Planning with theory of mind. *Trends in Cognitive Sciences*, 26(11): 959–971.
- Lakemeyer, G.; and Levesque, H. J. 2011. A semantic characterization of a useful fragment of the situation calculus with knowledge. *Artificial Intelligence*, 175(1): 142–164.
- Liu, D.; and Feng, Q. 2023. On the progression of belief. *Artificial Intelligence*, 322: 103947.
- Liu, Q.; and Liu, Y. 2018. Multi-agent Epistemic Planning with Common Knowledge. In *IJCAI*, 1912–1920.
- Löwe, B.; Pacuit, E.; and Witzel, A. 2011. DEL planning and some tractable cases. In *Logic, Rationality, and Interaction: Third International Workshop, LORI 2011, Guangzhou, China, October 10-13, 2011. Proceedings 3*, 179–192. Springer.
- McCarthy, J.; and Hayes, P. J. 1981. Some philosophical problems from the standpoint of artificial intelligence. In *Readings in artificial intelligence*, 431–450. Elsevier.
- Muise, C.; Belle, V.; Felli, P.; McIlraith, S.; Miller, T.; Pearce, A. R.; and Sonenberg, L. 2022. Efficient multi-agent epistemic planning: Teaching planners about nested belief. *Artificial Intelligence*, 302: 103605.
- Peloquin, C. 2023. *Representing Minds Representing Minds: An Examination of the Association between Recursive Mental State Attributions and Executive Processes*. Ph.D. thesis, Open Access Te Herenga Waka-Victoria University of Wellington.
- Pietarinen, A.-V. 2003. What do epistemic logic and cognitive science have to do with each other? *Cognitive systems research*, 4(3): 169–190.
- Reiter, R. 2001. *Knowledge in action: logical foundations for specifying and implementing dynamical systems*. MIT press.
- Thielscher, M. 2017. GDL-III: A Description Language for Epistemic General Game Playing. In *IJCAI*, 1276–1282.
- Wan, H.; Fang, B.; and Liu, Y. 2021. A general multi-agent epistemic planner based on higher-order belief change. *Artificial Intelligence*, 301: 103562.
- Ware, S. G.; and Siler, C. 2021a. Sabre: A narrative planner supporting intention and deep theory of mind. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 17, 99–106.
- Ware, S. G.; and Siler, C. 2021b. The sabre narrative planner: multi-agent coordination with intentions and beliefs. In *AAMAS Conference proceedings*.
- Wimmer, H.; and Perner, J. 1983. Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception. *Cognition*, 13(1): 103–128.
- Xiong, L.; and Liu, Y. 2016. Strategy Representation and Reasoning for Incomplete Information Concurrent Games in the Situation Calculus. In *IJCAI*, 1322–1329.

Yu, Q.; Wen, X.; and Liu, Y. 2013. Multi-Agent Epistemic Explanatory Diagnosis via Reasoning about Actions. In *IJ-CAI*, 1183–1190.