A GENTLE INTRODUCTION TO EPISTEMIC PLANNING

FOUNDATIONS AND CHALLENGES

Alessandro Burigana
Free University of Bozen-Bolzano

February 16th, 2024

Department of Mathematics,
Computer Science and Physics
University of Udine

A Picture from the Top

Epistemic Planning is the enrichment of classical automated planning, where the notions of knowledge and belief are introduced.

A Picture from the Top

Epistemic Planning is the enrichment of classical automated planning, where the notions of knowledge and belief are introduced.

Epistemic

- When does an agent knows or believe something?
- How do we represent the knowledge/ beliefs of multiple agents?

A Picture from the Top

Epistemic Planning is the enrichment of classical automated planning, where the notions of knowledge and belief are introduced.

Epistemic

- When does an agent knows or believe something?
- How do we represent the knowledge/ beliefs of multiple agents?

Planning

- How do we represent actions that change what agents know or believe?
- How do such actions change the current knowledge/beliefs of the agents?

A Picture from the Top

Epistemic Planning is the enrichment of classical automated planning, where the notions of knowledge and belief are introduced.

Epistemic

- When does an agent knows or believe something?
- How do we represent the knowledge/ beliefs of multiple agents?

$\Downarrow$

Epistemic Logic

Planning

- How do we represent actions that change what agents know or believe?
- How do such actions change the current knowledge/beliefs of the agents?

$\Downarrow$

Dynamic Epistemic Logic

A (SLIGHTLY) PHILOSOPHICAL INTRODUCTION

To Know or to Believe? . . .

What does it mean to believe that something is the case?

To Know or to Believe? . . .

What does it mean to believe that something is the case?

Belief

Belief is a **propositional attitude** that something is true.

→ **Mental state** held by an agent or organism toward a proposition.

To Know or to Believe? . . .

What does it mean to believe that something is the case?

Belief

Belief is a **propositional attitude** that something is true.

→ **Mental state** held by an agent or organism toward a proposition.

And what does it mean to know that something is the case?

To Know or to Believe? ...

What does it mean to believe that something is the case?

Belief

Belief is a **propositional attitude** that something is true.

$\rightarrow$ **Mental state** held by an agent or organism toward a proposition.

And what does it mean to know that something is the case?

The Tripartite Analysis of Knowledge

S knows that p iff

1. p is true;
2. S believes that p; and
3. S is justified in believing that p.

To Know or to Believe? ...

What does it mean to believe that something is the case?

Belief

Belief is a **propositional attitude** that something is true.

$\rightarrow$ **Mental state** held by an agent or organism toward a proposition.

And what does it mean to know that something is the case?

The Tripartite Analysis of Knowledge

S knows that p iff

1. p is true;
2. S believes that p; and
3. S is justified in believing that p.

$\Rightarrow$ **Justified True Belief (JTB)**

Is Justified True Belief actually Knowledge?

John is standing outside a field and, within it, he sees what looks exactly like a sheep.

→ Does John know that there is a sheep if the field?

Let's analyse the situation:

1. John sure believes that a sheep if the field.
2. John is also justified in believing so: he clearly sees it!
3. But **is it true** that there is a sheep in the field?

Is Justified True Belief actually Knowledge?

John is standing outside a field and, within it, he sees what looks exactly like a sheep.

→ Does John know that there is a sheep if the field?

Let's analyse the situation:

1. John sure believes that a sheep if the field.
2. John is also justified in believing so: he clearly sees it!
3. But **is it true** that there is a sheep in the field?

What John does **not** realize is that what he sees is actually a dog, disguised as a sheep.

→ Can we now say that now John knows that there is a sheep if the field?

Is Justified True Belief actually Knowledge?

John is standing outside a field and, within it, he sees what looks exactly like a sheep.

→ Does John know that there is a sheep if the field?

Let's analyse the situation:

1. John sure believes that a sheep if the field.
2. John is also justified in believing so: he clearly sees it!
3. But **is it true** that there is a sheep in the field?

What John does **not** realize is that what he sees is actually a dog, disguised as a sheep.

→ Can we now say that now John knows that there is a sheep if the field?

Moreover, there is actually a sheep behind the hill in the middle of the field.

→ What can we say now?

EPISTEMIC LOGIC

Syntax

Let $\mathcal{P}$ be a finite set of propositional atoms and $\mathcal{AG} = \{1, \ldots, n\}$ a finite set of agents. The **language $\mathcal{L}_{\mathcal{P},\mathcal{AG}}$ of Epistemic Logic** is given by the BNF:

Definition (Language of Epistemic Logic)

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \Box_i \varphi,$$

→ Operator $\Box_i$: depending on the context, describes what agent i knows or believes.

→ Dual operator $\Diamond_i$: describes what agent i considers to be possible or compatible.

Semantics

An *epistemic state* represents both factual information and what agents know/believe.

Definition (Epistemic Model)

An *epistemic model* is a triple $M = (W, R, V)$, where:

Semantics

An *epistemic state* represents both factual information and what agents know/believe.

Definition (Epistemic Model)

An *epistemic model* is a triple $M = (W, R, V)$, where:

- $W \neq \varnothing$ is a finite set of **possible worlds**;

Example

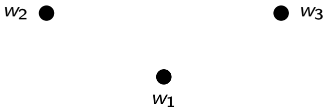

Semantics

An *epistemic state* represents both factual information and what agents know/believe.

Definition (Epistemic Model)

An *epistemic model* is a triple $M = (W, R, V)$, where:

- $W \neq \varnothing$ is a finite set of **possible worlds**;
- $R : \mathcal{AG} \to 2^{W \times W}$ assigns to each agent i an **accessibility relation** R_i;

Example

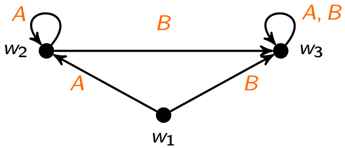

Semantics

An *epistemic state* represents both factual information and what agents know/believe.

Definition (Epistemic Model)

An *epistemic model* is a triple $M = (W, R, V)$, where:

- $W \neq \varnothing$ is a finite set of **possible worlds**;
- $R : \mathcal{AG} \rightarrow 2^{W \times W}$ assigns to each agent i an **accessibility relation** R_i;
- $V : \mathcal{P} \rightarrow 2^{W}$ is a **valuation function**; and

Example

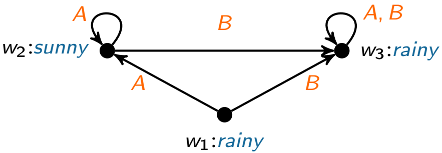

An *epistemic state* represents both factual information and what agents know/believe.

Definition (Epistemic Model)

An *epistemic model* is a triple $M = (W, R, V)$, where:

- $W \neq \varnothing$ is a finite set of **possible worlds**;
- $R : \mathcal{AG} \to 2^{W \times W}$ assigns to each agent i an **accessibility relation** R_i;
- $V : \mathcal{P} \to 2^W$ is a **valuation function**; and

Definition (Epistemic State)

An *epistemic state* is a pair (M, W_d) s.t. $W_d \subseteq W$ is a non-empty set of **designated worlds**.

Example

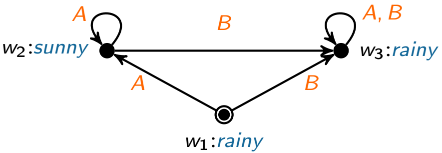

Semantics

Definition (Truth)

Let $s = (M, W_d)$, where $M = (W, R, V)$, be an *epistemic state* and let $w \in W$:

$$
\begin{array}{lll}
(M, w) \models p & \text{iff} & w \in V(p) \\
(M, w) \models \neg\varphi & \text{iff} & (M, w) \not\models \varphi \\
(M, w) \models \varphi \wedge \psi & \text{iff} & (M, w) \models \varphi \text{ and } (M, w) \models \psi \\
(M, w) \models \Box_i \varphi & \text{iff} & \forall v \text{ if } wR_i v \text{ then } (M, v) \models \varphi
\end{array}
$$

Moreover, $(M, W_d) \models \varphi$ iff $\forall w$ if $w \in W_d$ then $(M, w) \models \varphi$.

Example

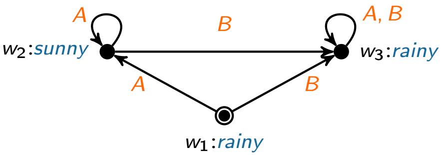

- $\Box_{Anne}\, sunny$
- $\Box_{Bob}\, rainy$
- $\Box_{Anne}\Box_{Bob}\, rainy$
- $\Diamond_{Bob}\Box_{Anne}\, rainy$

. . . That is the Question

How can epistemic states represent the knowledge and the beliefs of agents?

$\rightarrow$ We model them via **axioms**.

... That is the Question

How can epistemic states represent the knowledge and the beliefs of agents?

$\rightarrow$ We model them via **axioms**.

	Axiom	Frame Property	Knowledge	Belief
K	$\Box_i(\varphi \rightarrow \psi) \rightarrow (\Box_i\varphi \rightarrow \Box_i\psi)$	-	✓	✓
T	$\Box_i\varphi \rightarrow \varphi$	Reflexivity	✓	
D	$\Box_i\varphi \rightarrow \Diamond_i\varphi$	Seriality	✓	✓
4	$\Box_i\varphi \rightarrow \Box_i\Box_i\varphi$	Transitivity	✓	✓
5	$\neg\Box_i\varphi \rightarrow \Box_i\neg\Box_i\varphi$	Euclideanness	✓	✓

An epistemic state represents:

- Knowledge, when it satisfies axioms **K**, **T**, **4** and **5** $\Rightarrow$ **Logic S5**$_n$
- Belief, when it satisfies axioms **K**, **D**, **4** and **5** $\Rightarrow$ **Logic KD45**$_n$

DYNAMIC EPISTEMIC LOGIC

Actions in Classical Planning

Classical actions are:

1. **Propositional**
2. **Single-agent**
3. **Fully Observable**
4. **Deterministic**

Actions in Classical Planning

Example (Blocks World)

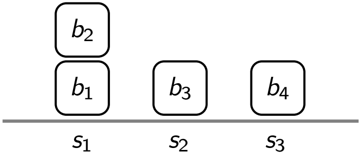

Classical actions are:

1. **Propositional**
2. **Single-agent**
3. **Fully Observable**
4. **Deterministic**

Actions in Classical Planning

Classical actions are:

1. **Propositional**
2. **Single-agent**
3. **Fully Observable**
4. **Deterministic**

Example (Blocks World)

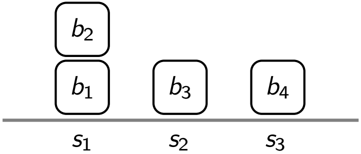

Action *move*(b, x, y):

- $\text{Pre}(move(b, x, y)) = On(b, x) \wedge Clear(b) \wedge Clear(y)$
- $\text{Eff}(move(b, x, y)) = \{On(b, y), Clear(x), \neg On(b, x), \neg Clear(y)\} \triangleright \top$

Classical actions are:

1. **Propositional**
2. **Single-agent**
3. **Fully Observable**
4. **Deterministic**

Example (Blocks World)

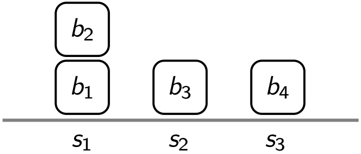

Action $move(b, x, y)$:

- $\text{Pre}(move(b, x, y)) = On(b, x) \land Clear(b) \land Clear(y)$
- $\text{Eff}(move(b, x, y)) = \{On(b, y), Clear(x), \neg On(b, x), \neg Clear(y)\} \triangleright \top$

$\rightarrow$ We now incrementally move from classical actions to epistemic actions.

Epistemic Blocks World

Example (Epistemic Blocks World)

- Agent a: only sees from above.

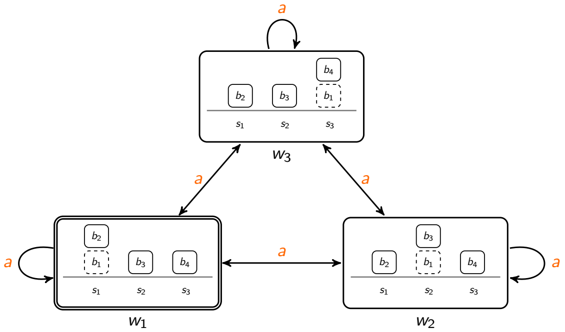

Epistemic Blocks World

Example (Multi-Agent Epistemic Blocks World)

- Agent a: only sees from above.
- Agent l: only sees from a top left position.

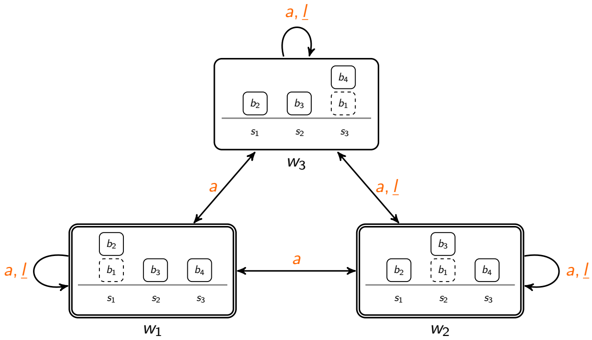

Epistemic Blocks World

Example (Multi-Agent Epistemic Blocks World)

- Agent a: only sees from above.
- Agent r: only sees from a top right position.

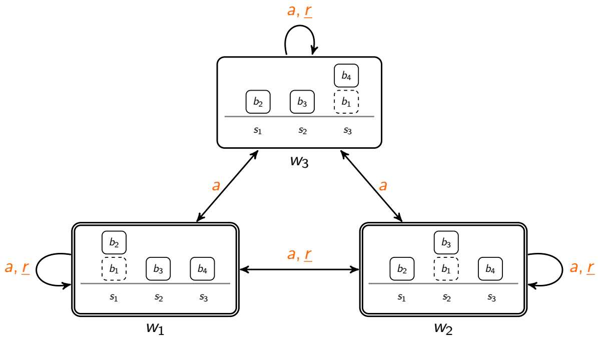

Epistemic Blocks World

Example (Multi-Agent Epistemic Blocks World)

- Agent a: only sees from above.
- Agent l: only sees from a top left position.
- Agent r: only sees from a top right position.

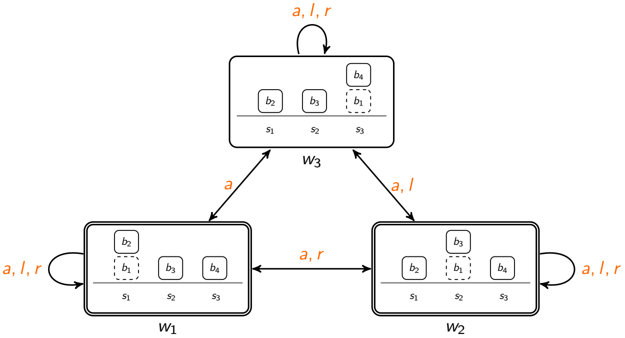

Epistemic Actions

Definition (Event Model)

An *event model* is a quadruple $\mathcal{E} = (E, Q, pre, post)$, where:

- $E \neq \varnothing$ is a finite set of **events**;
- $Q : \mathcal{AG} \to 2^{E \times E}$ assigns to each agent i an **accessibility relation** Q_i;

Intuitively:

- An event can be seen as a classical action.
- Accessibility relations specify the perspectives of agents on which events take place.

Epistemic Actions

Definition (Event Model)

An *event model* is a quadruple $\mathcal{E} = (E, Q, pre, post)$, where:

- $E \neq \varnothing$ is a finite set of **events**;
- $Q : \mathcal{AG} \rightarrow 2^{E \times E}$ assigns to each agent i an **accessibility relation** Q_i;
- $pre : E \rightarrow \mathcal{L}_{\mathcal{P},\mathcal{AG}}$ assigns to each event a **precondition**;
- $post : E \rightarrow (\mathcal{P} \rightarrow \mathcal{L}_{\mathcal{P},\mathcal{AG}})$ assigns to each event and atom a **postcondition**.

Intuitively:

- An event can be seen as a classical action, **each with its own pre- and postconditions**.
- Accessibility relations specify the perspectives of agents on which events take place.

Epistemic Actions

Definition (Event Model)

An *event model* is a quadruple $\mathcal{E} = (E, Q, pre, post)$, where:

- $E \neq \varnothing$ is a finite set of **events**;
- $Q : \mathcal{AG} \to 2^{E \times E}$ assigns to each agent i an **accessibility relation** Q_i;
- $pre : E \to \mathcal{L}_{\mathcal{P},\mathcal{AG}}$ assigns to each event a **precondition**;
- $post : E \to (\mathcal{P} \to \mathcal{L}_{\mathcal{P},\mathcal{AG}})$ assigns to each event and atom a **postcondition**.

Intuitively:

- An event can be seen as a classical action, **each with its own pre- and postconditions**.
- Accessibility relations specify the perspectives of agents on which events take place.

Definition (Epistemic Action)

An *epistemic action* is a pair $(\mathcal{E}, E_d)$, s.t. $E_d \subseteq E$ is a non-empty set of **designated events**.

Product Update

An action $(\mathcal{E}, E_d)$ is **applicable** is an epistemic state (M, W_d) iff for each designated world $w \in W_d$ there exists a designated event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition (Product Update)

Given (M, W_d) and $(\mathcal{E}, E_d)$, where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$, their **product update** $(M, W_d) \otimes (\mathcal{E}, E_d)$ is the epistemic state $((W', R', V'), W'_d)$ where:

Product Update

An action $(\mathcal{E}, E_d)$ is **applicable** is an epistemic state (M, W_d) iff for each designated world $w \in W_d$ there exists a designated event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition (Product Update)

Given (M, W_d) and $(\mathcal{E}, E_d)$, where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$, their **product update** $(M, W_d) \otimes (\mathcal{E}, E_d)$ is the epistemic state $((W', R', V'), W'_d)$ where:

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;

Product Update

An action $(\mathcal{E}, E_d)$ is **applicable** is an epistemic state (M, W_d) iff for each designated world $w \in W_d$ there exists a designated event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition (Product Update)

Given (M, W_d) and $(\mathcal{E}, E_d)$, where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$, their **product update** $(M, W_d) \otimes (\mathcal{E}, E_d)$ is the epistemic state $((W', R', V'), W'_d)$ where:

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;

Product Update

An action $(\mathcal{E}, E_d)$ is **applicable** is an epistemic state (M, W_d) iff for each designated world $w \in W_d$ there exists a designated event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition (Product Update)

Given (M, W_d) and $(\mathcal{E}, E_d)$, where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$, their **product update** $(M, W_d) \otimes (\mathcal{E}, E_d)$ is the epistemic state $((W', R', V'), W'_d)$ where:

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and

Product Update

An action $(\mathcal{E}, E_d)$ is **applicable** is an epistemic state (M, W_d) iff for each designated world $w \in W_d$ there exists a designated event $e \in E_d$ such that $(M, w) \models pre(e)$.

Definition (Product Update)

Given (M, W_d) and $(\mathcal{E}, E_d)$, where $M = (W, R, V)$ and $\mathcal{E} = (E, Q, pre, post)$, their **product update** $(M, W_d) \otimes (\mathcal{E}, E_d)$ is the epistemic state $((W', R', V'), W'_d)$ where:

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example

Public Announcement

Agent r tells everybody that he knows that $\neg On(b_1, s_3)$.

$e : \langle \Box_r \neg On(b_1, s_3), \top \rangle$

Example

Public Announcement

Agent r tells everybody that he knows that $\neg On(b_1, s_3)$.

$e : \langle \Box_r \neg On(b_1, s_3), \top \rangle$

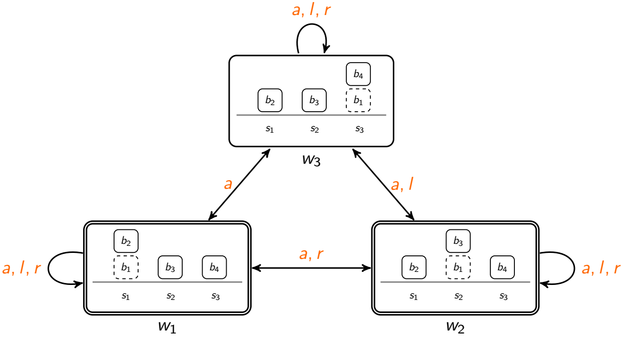

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example

Public Announcement

Agent r tells everybody that he knows that $\neg On(b_1, s_3)$.

$e : \langle \Box_r \neg On(b_1, s_3), \top \rangle$

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_iv \text{ and } eQ_if\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

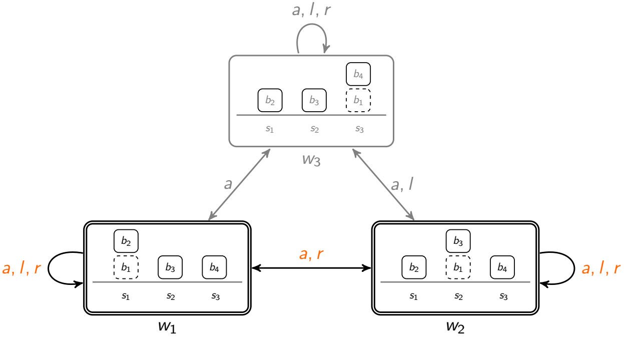

Public Announcements

Example

Public Announcement

Agent r tells everybody that he knows that $\neg On(b_1, s_3)$.

$e : \langle \Box_r \neg On(b_1, s_3), \top \rangle$

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

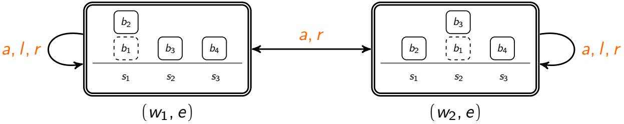

Semi-Private Sensing Action

Example

Semi-Private Sensing Action

Agent r peeks under block b_2 while agents a and l observe him. Specifically:

- Agents r and l observe what is actually being sensed.
- Agent a can not directly observe what agent r is seeing.

Trivial postconditions are omitted.

Semi-Private Sensing Action

Example

Semi-Private Sensing Action

Agent r peeks under block b_2 while agents a and l observe him. Specifically:

- Agents r and l observe what is actually being sensed.
- Agent a can not directly observe what agent r is seeing.

Trivial postconditions are omitted.

Definition (Product Update)

- $W' = \{(w,e) \in W \times E \mid (M,w) \models pre(e)\}$;
- $R'_i = \{((w,e),(v,f)) \in W' \times W' \mid wR_iv \text{ and } eQ_if\}$;
- $V'(p) = \{(w,e) \in W' \mid (M,w) \models post(e)(p)\}$; and
- $W'_d = \{(w,e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

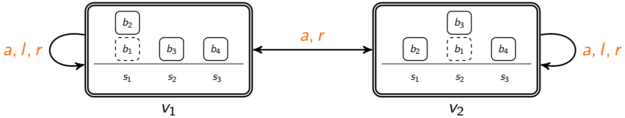

Example

Semi-Private Sensing Action

Agent r peeks under block b_2 while agents a and l observe him. Specifically:

- Agents r and l observe what is actually being sensed.
- Agent a can not directly observe what agent r is seeing.

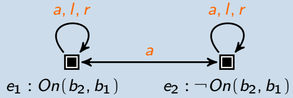

Trivial postconditions are omitted.

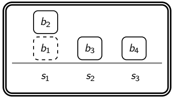

(v_1, e_1)

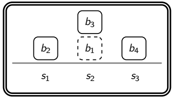

(v_2, e_2)

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v$ and $eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d$ and $e \in E_d\}$.

Semi-Private Sensing Action

Example

Semi-Private Sensing Action

Agent r peeks under block b_2 while agents a and l observe him. Specifically:

- Agents r and l observe what is actually being sensed.
- Agent a can not directly observe what agent r is seeing.

Trivial postconditions are omitted.

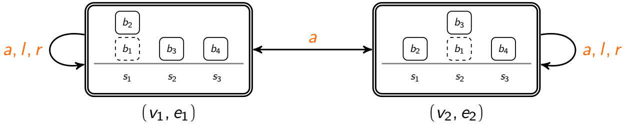

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R_i' = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W_d' = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example

Public Sensing Action

All agents peek under block b_2.

→ **Non-deterministic** action!

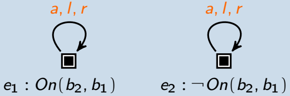

Trivial postconditions are omitted.

Example

Public Sensing Action

All agents peek under block b_2.

$\rightarrow$ **Non-deterministic** action!

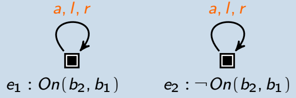

Trivial postconditions are omitted.

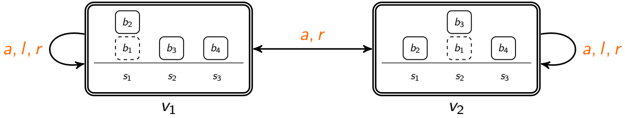

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example

Public Sensing Action

All agents peek under block b_2.

→ **Non-deterministic** action!

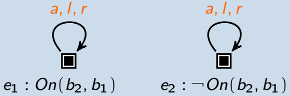

Trivial postconditions are omitted.

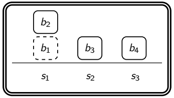

(v_1, e_1)

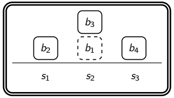

(v_2, e_2)

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Public Sensing Action

Example

Public Sensing Action

All agents peek under block b_2.

$\rightarrow$ **Non-deterministic** action!

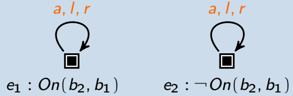

Trivial postconditions are omitted.

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

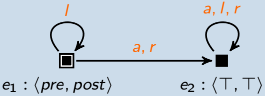

Private Ontic Actions

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

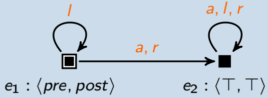

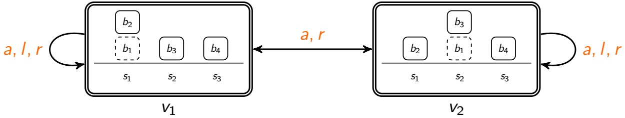

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Private Ontic Actions

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

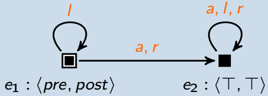

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

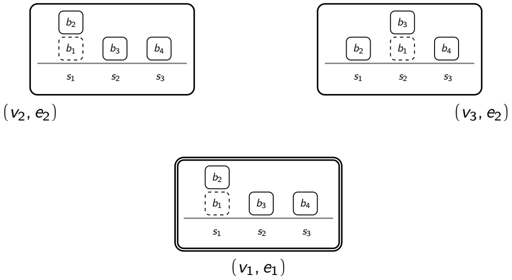

Private Ontic Actions

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

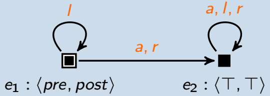

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

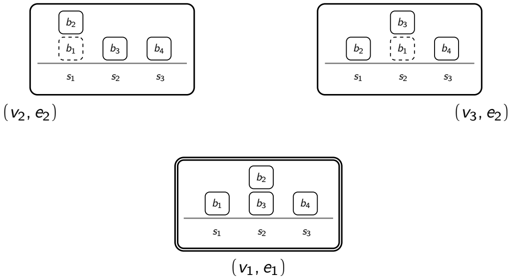

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

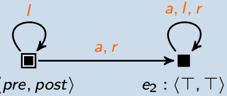

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

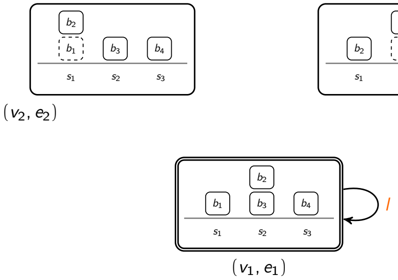

Example

Private Ontic Action

Agent l privately moves block b_2 from b_1 to b_3, where:

- $pre = On(b_2, b_1) \wedge Clear(b_2) \wedge Clear(b_3)$
- $post(e_1)(On(b_2, b_1)) = \bot$
- $post(e_1)(On(b_2, b_3)) = \top$

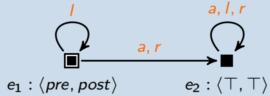

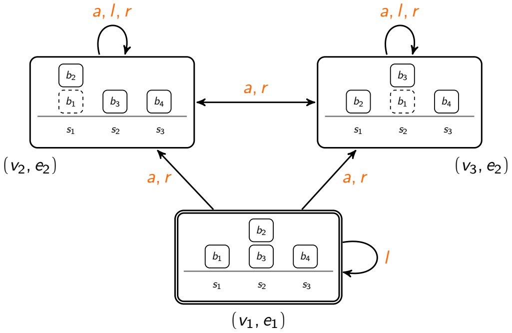

Definition (Product Update)

- $W' = \{(w, e) \in W \times E \mid (M, w) \models pre(e)\}$;
- $R'_i = \{((w, e), (v, f)) \in W' \times W' \mid wR_i v \text{ and } eQ_i f\}$;
- $V'(p) = \{(w, e) \in W' \mid (M, w) \models post(e)(p)\}$; and
- $W'_d = \{(w, e) \in W' \mid w \in W_d \text{ and } e \in E_d\}$.

Classical Vs. Epistemic Actions

To summarize:

Classical actions are:

1. **Propositional**
2. **Single-agent**
3. **Fully Observable**
4. **Deterministic**

Epistemic actions are:

1. **Modal**
2. **Multi-agent**
3. **Partially Observable**
4. **Non-deterministic**

Moreover, epistemic actions model both factual and higher-order knowledge change.

$\rightarrow$ There are **no restrictions** on the reasoning power of agents! (More on this later)

Epistemic Planning Task

Definition (Planning Task)

An **(epistemic) planning task** is a triple $T = (s_0, \mathcal{A}, \varphi_g)$, where:

- s_0 is an **initial epistemic state**;
- $\mathcal{A}$ is a finite **set of actions**;
- $\varphi_g \in \mathcal{L}_{\mathcal{P},\mathcal{AG}}$ is a **goal formula**.

Epistemic Planning Task

Definition (Planning Task)

An **(epistemic) planning task** is a triple $T = (s_0, \mathcal{A}, \varphi_g)$, where:

- s_0 is an **initial epistemic state**;
- $\mathcal{A}$ is a finite **set of actions**;
- $\varphi_g \in \mathcal{L}_{\mathcal{P},\mathcal{AG}}$ is a **goal formula**.

Definition (Solution)

A **solution** to a planning task $(s_0, \mathcal{A}, \varphi_g)$ is a finite sequence $\alpha_1, \ldots, \alpha_m$ of actions of $\mathcal{A}$ s.t.:

1. For each $1 \leqslant k \leqslant m$, α_k is applicable in $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_{k-1}$, and
2. $s_0 \otimes \alpha_1 \otimes \cdots \otimes \alpha_m \models \varphi_g$.

Epistemic Plan Existence Problem

Definition (Plan Existence Problem)

Let $n \geqslant 1$ and $\mathcal{T}$ be a class of planning tasks. PlanEx($\mathcal{T}$, n) is the following decision problem: "Given a planning task $T = (s_0, \mathcal{A}, \varphi_g) \in \mathcal{T}$, where $|\mathcal{AG}| = n$, does T have a solution?"

Theorem (Bolander and Andersen [BA11])

Let $\mathcal{T}$ be the class of all epistemic planning tasks and let $n \geqslant 1$. Then, PlanEx*($\mathcal{T}$, n) is* ***undecidable****.*

CURRENT CHALLENGES

Decidable Fragments

A great deal of effort has been spent over the past decade to devise **decidable** fragments of the epistemic plan existence problem.

Let $\mathcal{T}(a, b)$ denote the class of epistemic planning tasks where:

- a is the maximum modal depth of **preconditions**, and
- b is the maximum modal depth of **postconditions**. We indicate with $b = -1$ the absence of postconditions.

PlanEx($\mathcal{T}(0, -1)$, n)	PSPACE-complete [CMS16]
PlanEx($\mathcal{T}(1, -1)$, n)	Unknown [CMS16]
PlanEx($\mathcal{T}(2, -1)$, n)	UNDECIDABLE [CMS16]
PlanEx($\mathcal{T}(0, 0)$, n)	DECIDABLE [YWL13; AMP14]
PlanEx($\mathcal{T}(1, 0)$, n)	DECIDABLE [Bol+20]

Decidable Fragments (cont.)

Others have focused on considering the plan existence problem of tasks under well-known modal logics (Aucher and Bolander [AB13]).

<table>
<tr><th>Logic</th><th>Single-agent</th><th>Multi-agent</th></tr>
<tr><td>K</td><td rowspan="3">UNDECIDABLE</td><td rowspan="6">UNDECIDABLE</td></tr>
<tr><td>KT</td></tr>
<tr><td>K4</td></tr>
<tr><td>K45</td><td>DECIDABLE</td></tr>
<tr><td>S4</td><td>UNDECIDABLE</td></tr>
<tr><td>S5</td><td>DECIDABLE</td></tr>
</table>

A Semantic Approach

What if we combined the two previous approaches together?

$\rightarrow$ We can limit the reasoning power of agents via **modal axioms**.

A Semantic Approach

What if we combined the two previous approaches together?

$\rightarrow$ We can limit the reasoning power of agents via **modal axioms**.

Knowledge Commutativity

$$\textbf{C} \quad \Box_i \Box_j \varphi \rightarrow \Box_j \Box_i \varphi$$

We call **C-S5**$_n$ the logic S5$_n$ augmented with axiom **C**.

A Semantic Approach

What if we combined the two previous approaches together?

$\rightarrow$ We can limit the reasoning power of agents via **modal axioms**.

Knowledge Commutativity

$$\textbf{C} \quad \Box_i \Box_j \varphi \rightarrow \Box_j \Box_i \varphi$$

We call **C-S5$_n$** the logic S5$_n$ augmented with axiom **C**.

Lemma (Burigana et al. [Bur+23])

Let (M, W_d) be a bisimulation-contracted C-S5$_n$-state, with $M = (W, R, V)$. Then, $|W|$ ***is bounded in n and $|\mathcal{P}|$.***

Theorem (Burigana et al. [Bur+23])

The plan existence problem in ***C-S5$_n$*** *is* ***decidable****.*

Generalizing Commutativity

Let $b > 1$ be a fixed integer constant:

b-Commutativity

$$\mathbf{C}^b \quad (\Box_i \Box_j)^b \varphi \rightarrow (\Box_j \Box_i)^b \varphi$$

Let $1 < \ell \leqslant n$ be a fixed integer constant, let $\langle i_1, \ldots, i_\ell \rangle$ be a repetition-free sequence of agents and let π be any of its permutations:

Weak Commutativity

$$\mathbf{wC}_\ell \quad \Box_{i_1} \ldots \Box_{i_\ell} \varphi \rightarrow \Box_{\pi_{i_1}} \ldots \Box_{\pi_{i_\ell}} \varphi$$

- We call **$\mathbf{C}^b$-S5$_n$** the logic S5$_n$ augmented with axiom $\mathbf{C}^b$.
- We call **$\mathbf{wC}_\ell$-S5$_n$** the logic S5$_n$ augmented with axiom $\mathbf{wC}_\ell$ (for all π).

Benefits of Semantic Approach

We obtain positive results:

Logic	Decidability
K_n, K_n, KT_n, $K4_n$, $K45_n$, $S4_n$, $S5_n$	UNDECIDABLE [AB13]
C^b-$S5_n$ (n>2)	UNDECIDABLE [Bur+23]
C^b-$S5_2$	DECIDABLE [Bur+23]
wC_ℓ-$S5_n$	
C-$S5_n$	

- Well-known epistemic planning formalism are captured by C-$S5_n$.
- Flexible approach: different axioms can be devised depending on the situation.
- No strong restrictions on modal depth.

Efficient Implementations

Current approaches:

- Compilation of fragments of DEL into classical planning.
- Bisimulation-contraction techniques.
- Ad hoc implementations of fragments of DEL.

Efficient Implementations

Current approaches:

- Compilation of fragments of DEL into classical planning.
- Bisimulation-contraction techniques.
- Ad hoc implementations of fragments of DEL.

Future directions:

- **Symbolic** approaches: SMT encodings, syntactic models.
- Heuristics, heuristics, heuristics.
  $\rightarrow$ Currently working on: **Epistemic Planning Graph**.
- **Bounded** bisimulation contractions.

Benchmarks for Epistemic Planning

The many fragments of DEL are hard to compare:

→ Different ad hoc languages (if any) capture only a part of DEL.

→ Different custom benchmarks.

Benchmarks for Epistemic Planning

The many fragments of DEL are hard to compare:

→ Different ad hoc languages (if any) capture only a part of DEL.

→ Different custom benchmarks.

We need a unified language for the entire DEL semantics. This would allow the following:

- Standard language to represent epistemic planning domains.
- Development of a publicly available and shared set of benchmarks.
- Easier comparison of results.
  - → Better overall progress of efficient techniques.

Benchmarks for Epistemic Planning

The many fragments of DEL are hard to compare:

→ Different ad hoc languages (if any) capture only a part of DEL.

→ Different custom benchmarks.

We need a unified language for the entire DEL semantics. This would allow the following:

- Standard language to represent epistemic planning domains.
- Development of a publicly available and shared set of benchmarks.
- Easier comparison of results.
  - → Better overall progress of efficient techniques.

Currently working on **EPDDL**:

→ Borrows the well-known syntax of PDDL and extends it to capture the whole DEL semantics.

Belief Revision in DEL

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

Belief Revision in DEL

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

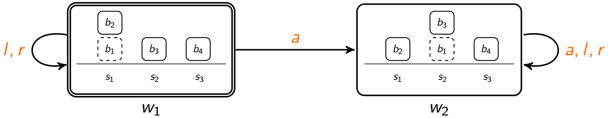

Belief Revision in DEL

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

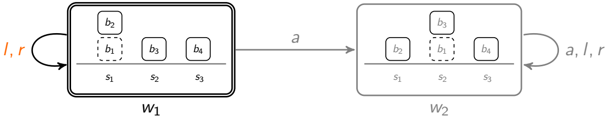

Belief Revision in DEL

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

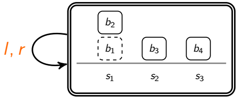

(w_1, e)

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

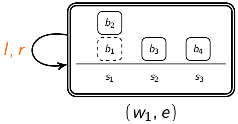

Some axioms might "break" after product update.

→ The state is no longer **serial** → Axiom **D** is not preserved.

→ The state does not represent what agent a believes.

Public Announcement

$e : \langle On(b_1, s_3), \top \rangle$

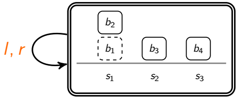

(w_1, e)

Some axioms might "break" after product update.

→ The state is no longer **serial** → Axiom **D** is not preserved.

→ The state does not represent what agent a believes.

How do we fix this?

- **Plausibility models**: belief of the agent is captured by the most plausible worlds.
  - → We recover **a**'s beliefs by looking at what he considers to be plausible.
- **Recovery**: prior to public announcements, we do a recovering action that "expands" the agents' beliefs.
- **Modifying the product update operator**.

Epistemic (outside of) Planning

Epistemic planning is still a relatively recent research area.

→ Many things still to address.

→ Has not been exploited in real scenarios.

Different areas would benefit from epistemic planning and reasoning:

- **Multi-Agent Systems**
  - → Self-driving vehicles
  - → Social commitments
  - → Business Process Management
- **Legal reasoning**
- Virtually any scenario involving uncertainty and/or different perspectives

THANK YOU

Questions?