

Dynamic epistemic logic

Tristan Charrier François Schwarzentruber



École Normale Supérieure Rennes

May 13, 2019

Outline

- 1 Discussion about modeling actions
 - In the verification/model checking community
 - In philosophy / AI
 - Syntactic specifications
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Outline

- 1 Discussion about modeling actions
 - In the verification/model checking community
 - In philosophy / AI
 - Syntactic specifications
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

In the verification/model checking community

```
// HelloWin.cpp : Defines the entry point for the application.
//
#include "stdafx.h"
#include "resource.h"
#define MAX_LOADSTRING 100

// Global Variables
HINSTANCE hInst;                                // current instance
TCHAR szTitle[MAX_LOADSTRING];                 // The title b
TCHAR szWindowClass[MAX_LOADSTRING];          // The

// Forward declarations of functions included in this code module:
ATOM                MyRegisterClass(HINSTANCE hInstance);
BOOL                InitInstance(HINSTANCE, int);
LRESULT CALLBACK    WndProc(HWND, UINT, WPARAM, LPARAM);
LRESULT CALLBACK    About(HWND, UINT, WPARAM, LPARAM);



int APIENTRY WinMain(HINSTANCE hInstance,
                    HINSTANCE hPrevInstance,
                    LPSTR lpCmdLine,
                    int nCmdShow)
{
    // TODO: Place code here.
    MSG msg;
    HACCEL hAccelTable;

    // Initialize global strings
    LoadString(hInstance, IDS_APP_TITLE, szTitle, MAX_LOADSTRING);
    LoadString(hInstance, IDC_HELLOWIN, szWindowClass, MAX_LOADSTRING);
    MyRegisterClass(hInstance);

    // Perform application initialization:
    if (!InitInstance (hInstance, nCmdShow))
    {

```

Program





Transition system

Action = an edge \longrightarrow

In the verification/model checking community

```
// HelloWin.cpp : Defines the entry point for the application.
//
#include "stdafx.h"
#include "resource.h"
#define MAX_LOADSTRING 100
// Global Variables
HINSTANCE hInst          // current instance
TCHAR szTitle[MAX_LOADSTRING] // The title b
TCHAR szWindowClass[MAX_LOADSTRING]; // The
// Forward declarations of functions included in this code module:
ATOM            MyRegisterClass(HINSTANCE hInstance);
BOOL            InitInstance(HINSTANCE, int);
LRESULT CALLBACK WndProc(HWND, UINT, WPARAM, LPARAM);
LRESULT CALLBACK About(HWND, UINT, WPARAM, LPARAM);
int APIENTRY WinMain(HINSTANCE hInstance,
                    HINSTANCE hPrevInstance,
                    LPSTR lpCmdLine,
                    int nCmdShow)
{
    // TODO: Place code here.
    MSG msg;
    HACCEL hAccelTable;
    // Initialize global strings
    LoadString(hInstance, IDS_APP_TITLE, szTitle, MAX_LOADSTRING);
    LoadString(hInstance, IDC_HELLOVIN, szWindowClass, MAX_LOADSTRING);
    MyRegisterClass(hInstance);
    // Perform application initialization:
    if (!InitInstance (hInstance, nCmdShow))
    {
```

Program



Transition system

Action = an edge \longrightarrow

Epistemic = edges — —

Outline

- 1 Discussion about modeling actions
 - In the verification/model checking community
 - **In philosophy / AI**
 - Syntactic specifications
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

In philosophy / AI

The mechanism of actions is important.

Public/private announcement Announce ‘She knows you hold $5\heartsuit$ ’

Public action play card $5\heartsuit$

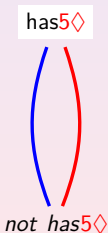

Private action secretly remove card $5\heartsuit$

Belief revision learn p although believing $\neg p$

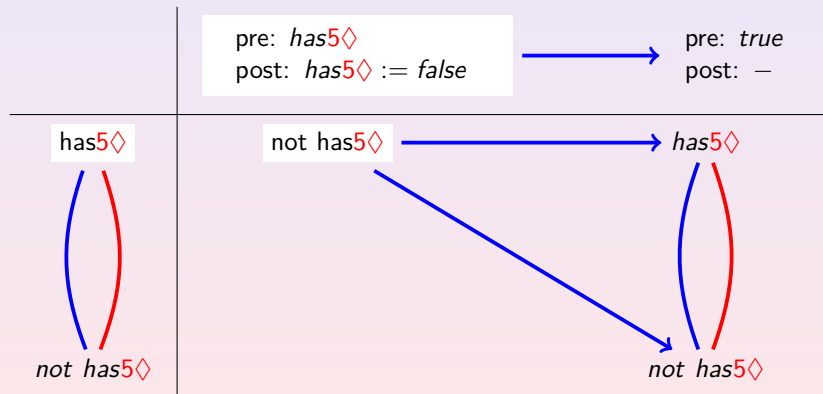
Solution: Dynamic epistemic logic

	State	Action
Classical planning	has5◇	pre: <i>has5◇</i> post: <i>has5◇ := false</i>

Solution: Dynamic epistemic logic

	State	Action
Classical planning	has5◇	pre: has5◇ post: has5◇ := false
DEL [Baltag et al. TARK 1998] [van Ditmarsch et al. 2007] = Kripkean models of classical planning	 <p>has5◇</p> <p>not has5◇</p>	pre: has5◇ post: has5◇ := false  <p>pre: true post: —</p>








Computing the next state: product update



Outline

- 1 Discussion about modeling actions
 - In the verification/model checking community
 - In philosophy / AI
 - Syntactic specifications
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Syntactic specifications



Game description language [Love et al. 2008] [Thielscher, IJCAI 2017]	agent a sees the game position
Flatland  [Balbiani et al., IGPL 2014]  [Gasquet, Goranko, __, AAMAS 2014]  [Gasquet, Goranko, __, JAAMAS 2016]	agent a sees agent b
Visibility atoms  [Charrier et al. KR 2016]	a sees the truth value of p
Paying attention to public announcements  [Bolander et al. JoLLI 2016]	$B_a \text{payAtt}(b) \rightarrow [p!] B_a B_b p$
Asynchronous announcements  [Knight et al. MS in CS 2019]	$[p!][read_a] K_a p$
Epistemic gossip  [van Ditmarsch et al., JAL 2017]	$[call_{ab}] K_a secret_b$

From DEL to epistemic temporal logics

Syntactic specification

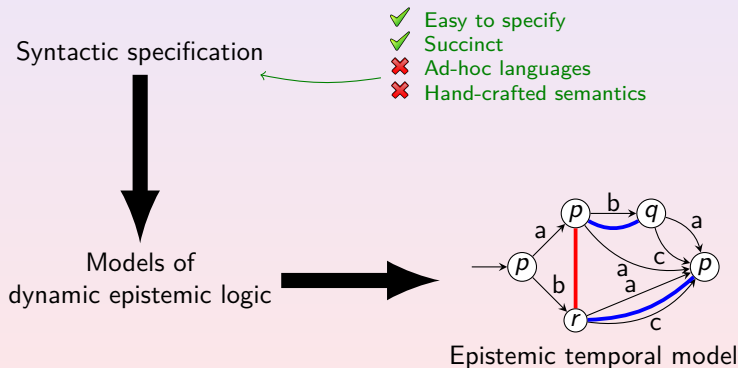


Models of
dynamic epistemic logic

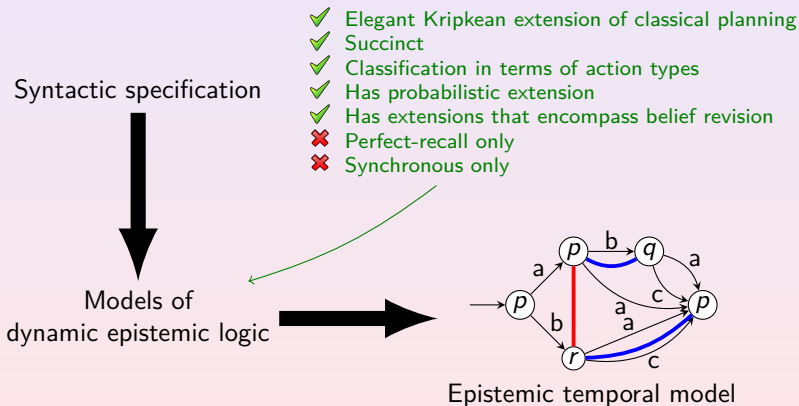


Epistemic temporal model

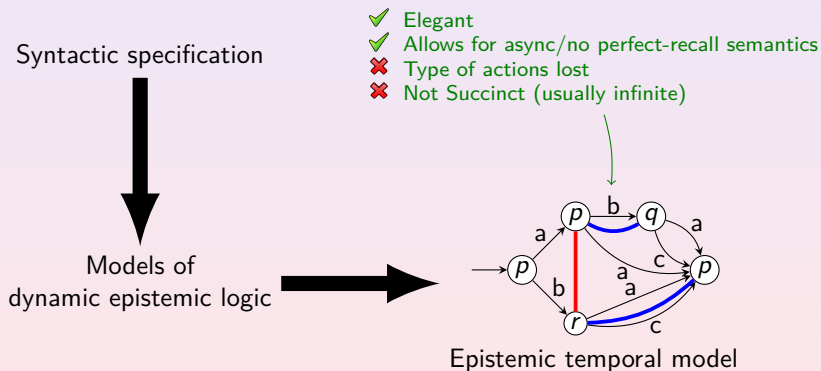
From DEL to epistemic temporal logics



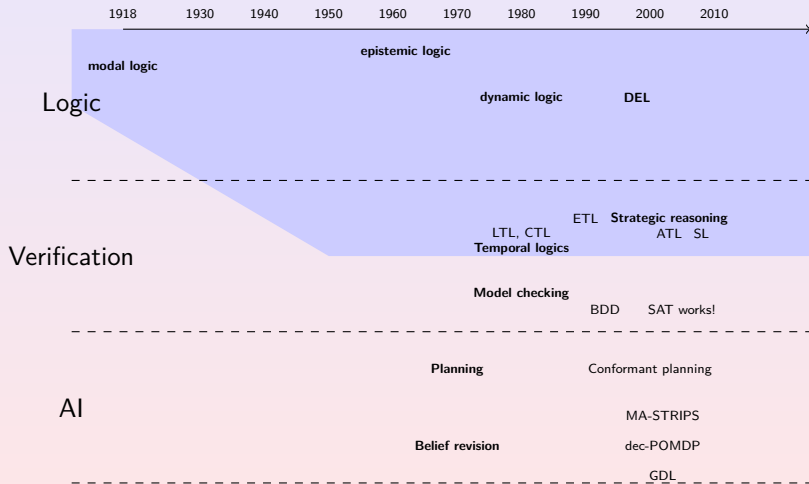
From DEL to epistemic temporal logics



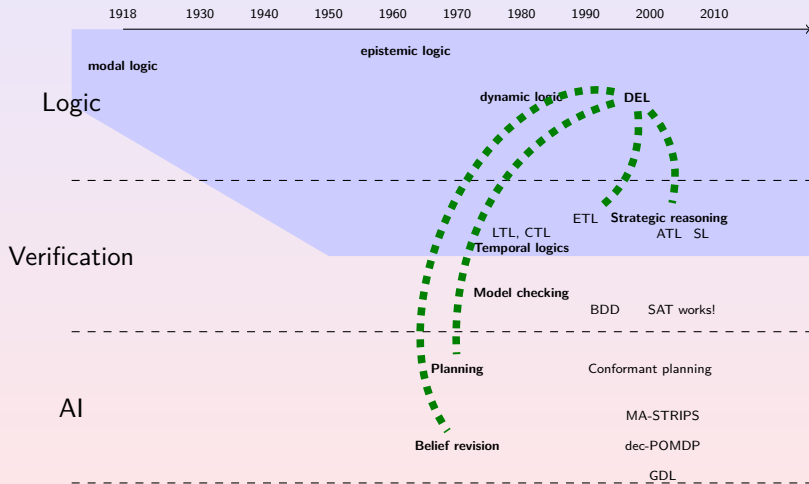
From DEL to epistemic temporal logics



Timeline



Timeline



Outline

- 1 Discussion about modeling actions
- 2 **Formal definition of event models**
 - Examples of actions
 - Definition
 - Effect of actions
 - Dynamic language
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
 - Examples of actions
 - Definition
 - Effect of actions
 - Dynamic language
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Examples of actions

[baltag1998logic]

Example (Public announcement of "p")


→ pre: p
 post: $-$ ↻ a, b

Example (Private announcement "p" to a)


→ pre: p
 post: $-$
 ↻ a

→ b → pre: **true**
 post: $-$ ↻ a, b

Examples of actions




Example (Transfer marble from basket to box)



Outline

- 1 Discussion about modeling actions
- 2 **Formal definition of event models**
 - Examples of actions
 - **Definition**
 - Effect of actions
 - Dynamic language
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Actions



Definition

An **event model** $\mathcal{E} = (E, (R_a^{\mathcal{E}})_{a \in AGT}, pre, post)$ is a tuple where:


- $E = \{e, e', \dots\}$ is a non-empty finite set of possible **events**,
- $R_a^{\mathcal{E}} \subseteq E \times E$ is an **accessibility relation** on E for agent a ,
- $pre : E \rightarrow \mathcal{L}_{EL}$ is a **precondition function**,
- $post : E \times AP \rightarrow \mathcal{L}_{EL}$ is a **postcondition function**.

A pair (\mathcal{E}, e) is called an **action**, where e represents the actual event of (\mathcal{E}, e) .

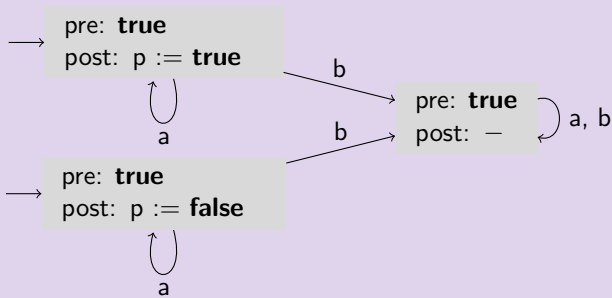
A pair (\mathcal{E}, E_0) , for $E_0 \subseteq E$, is a **non-deterministic action**. The set E_0 is the set of **triggerable events**.

Deterministic and non-deterministic actions

Deterministic action = single-pointed event model (\mathcal{E}, e)



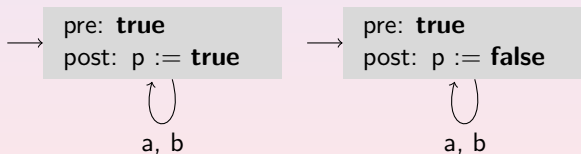
Non-deterministic action = multi-pointed event model



Public actions

Definition


An action is said to be *public* if the accessibility relations in underlying event model are self-loops.



Non-ontic actions

Definition

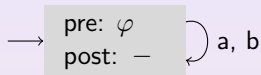
An action is said to be *non-ontic* if the postconditions are trivial: for all $e \in E$, for all propositions $p \in AP$, $post(e, p) = p$.




Outline

- 1 Discussion about modeling actions
- 2 **Formal definition of event models**
 - Examples of actions
 - Definition
 - **Effect of actions**
 - Dynamic language
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning




Effect of a public announcement of φ : only keep φ -worlds



In Hintikka's World: Try on several examples!



Example of an update product



Update product: formal definition

Let $\mathcal{M} = (W, \{R_a\}_{a \in AGT}, V)$ be an epistemic model and $\mathcal{E} = (E, (R_a^{\mathcal{E}})_{a \in AGT}, pre, post)$ be an event model.

Definition

The **update product** of \mathcal{M} and \mathcal{E} is the epistemic model $\mathcal{M} \otimes \mathcal{E} = (W^{\otimes}, \{R_a^{\otimes}\}_{a \in AGT}, V^{\otimes})$ where:

$$W^{\otimes} = \{(w, e) \in W \times E \mid \mathcal{M}, w \models pre(e)\},$$

$$R_a^{\otimes}(w, e) = \{(w', e') \in W^{\otimes} \mid wR_a w' \text{ and } eR_a^{\mathcal{E}} e'\},$$

$$V^{\otimes}(w, e) = \{p \in AP \mid \mathcal{M}, w \models post(e)(p)\}$$

Pointed update products

Definition

The **successor state** of an epistemic state (\mathcal{M}, w) by action (\mathcal{E}, e) is

$$(\mathcal{M}, w) \otimes (\mathcal{E}, e) =^{\text{def}} (\mathcal{M} \otimes \mathcal{E}, (w, e))$$

if $\mathcal{M}, w \models \text{pre}(e)$, otherwise it is undefined.

Notation

- We write e instead of (\mathcal{E}, e) ;
- We write the word ' we ' instead of the pair (w, e) ;
- We write $\mathcal{M} \otimes \mathcal{E}^n$ for $\mathcal{M} \otimes \mathcal{E} \otimes \dots \otimes \mathcal{E}$, n times.
- We write $we_1 \dots e_n \models \varphi$ instead of $\mathcal{M} \otimes \mathcal{E}^n, we_1 \dots e_n \models \varphi$.

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
 - Examples of actions
 - Definition
 - Effect of actions
 - **Dynamic language**
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Dynamic language

Definition

The language $\mathcal{L}_{\text{DELCK}}$ extends $\mathcal{L}_{\text{ELCK}}$ with dynamic modalities and is defined by the following BNF:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \vee \varphi) \mid K_a\varphi \mid C_G\varphi \mid \langle \mathcal{E}, E_0 \rangle \varphi$$

where \mathcal{E}, E_0 ranges over the set of non-deterministic actions.

Definition

We extend the definition $\mathcal{M}, w \models \varphi$ to $\mathcal{L}_{\text{DELCK}}$ with the following clause:

- $\mathcal{M}, w \models \langle \mathcal{E}, E_0 \rangle \varphi$ if there exists $e \in E_0$ s.th.
 $\mathcal{M}, w \models \text{pre}(e)$ and $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi$.

Dual operator

We define $[\mathcal{E}, E_0]$ to be $\neg\langle\mathcal{E}, E_0\rangle\neg$.

The semantics is:

- $\mathcal{M}, w \models [\mathcal{E}, E_0]\varphi$ if for all $e \in E_0$ we have
 $\mathcal{M}, w \models pre(e)$ implies $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi$;
- $\mathcal{M}, w \models \langle\mathcal{E}, E_0\rangle\varphi$ if there exists $e \in E_0$ s.th.
 $\mathcal{M}, w \models pre(e)$ and $\mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi$.

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
 - Examples of actions
 - Definition
 - Effect of actions
 - Dynamic language
 - Expressivity
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning

Expressivity

Theorem

DEL and **EL** have the same expressivity.

Idea: we remove the dynamic operators $[\mathcal{E}, E]$. Let us explain it just with public announcements:

$[\varphi!] \psi$: if φ holds then after having announced φ publicly, ψ holds.

$[\varphi!] p$	says the same thing than	$(\varphi \rightarrow p)$
$[\varphi!](\psi \wedge \chi)$	says the same thing than	$([\varphi!] \psi \wedge [\varphi!] \chi)$
$[\varphi!] \neg \psi$	says the same thing than	$(\varphi \rightarrow \neg [\varphi!] \psi)$
$[\varphi!] K_a \psi$	says the same thing than	$(\varphi \rightarrow K_a [\varphi!] \psi)$
$[\varphi!][\psi!] \chi$	says the same thing than	$[\varphi \wedge [\varphi!] \psi!] \chi$

General proof in [Baltag, Moss and Solecki, 2003a]

DEL is more succinct: [Lutz, AAMAS 2006]

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking**
 - Model checking problem
 - Complexity
- 4 Theorem proving
- 5 Epistemic planning

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking**
 - Model checking problem
 - Complexity
- 4 Theorem proving
- 5 Epistemic planning

Model checking problem

Definition (model checking problem)

- Input:

- An epistemic state



- A formula, e.g. $\langle action_1; action_2 \rangle K_a p$;

- Output: yes if




no otherwise.


Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking**
 - Model checking problem
 - Complexity**
- 4 Theorem proving
- 5 Epistemic planning

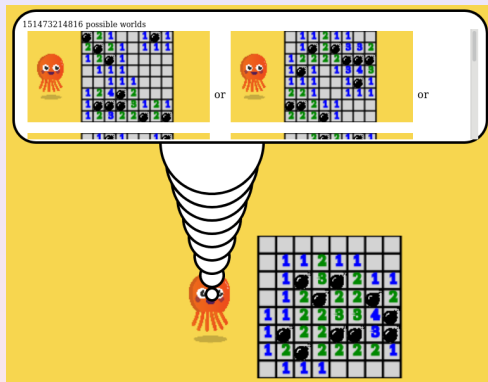
Model checking complexity



P-complete
[van Benthem, 2011]

PSPACE-complete
 [Aucher, __, TARK 2013]
[Pol et al. 2016]


State explosion problem



Example

Minesweeper easy 8×8 with 10 bombs: $> 10^{12}$ possible worlds.


State explosion problem





Example

Minesweeper 10×12 with 20 bombs: $> 10^{25}$ possible worlds.

Solution to the state explosion problem




[[DBLP:conf/lori/BenthemEGS15](#)], [[DBLP:journals/logcom/BenthemEGS18](#)]

 [Charrier _ AAMAS 2017],  [Charrier _ AiML 2018]

- Succinct representations of epistemic states **and** actions;
- Easy to specify by means of accessibility programs;
- Succinct model checking still in PSPACE.

Impact



Theoretical

Existence of a (uniform) strategy in **bounded** imperfect info games is in PSPACE.

Implementation: PSPACE techniques

Symbolic Model checking implemented in Hintikka's World:

- by Sébastien Gamblin and Alexandre Niveau (univ. Caen)
- using BDDs (C wrapper of CUDD compiled in wasm).

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving**
- 5 Epistemic planning

Theorem proving

Motivation: parametrized verification

for all epistemic states in which p holds:

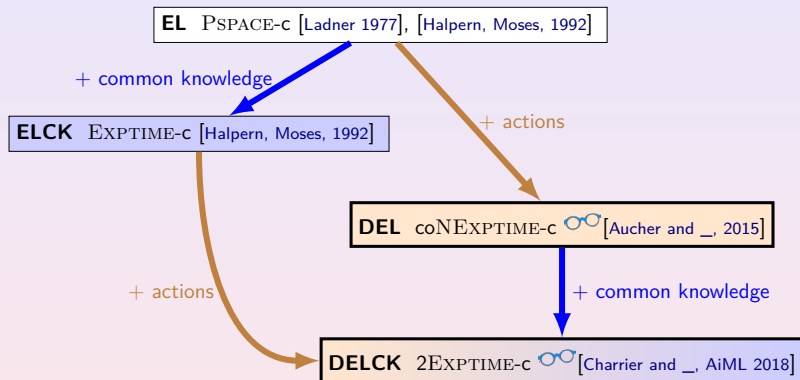


$p \rightarrow \langle action_1; action_2 \rangle K_a p$ is a *theorem*, i.e. true in all epistemic states.

Definition

- Input: a formula φ ;
- Output: yes if φ is a theorem, no otherwise.

Theorem proving is highly intractable



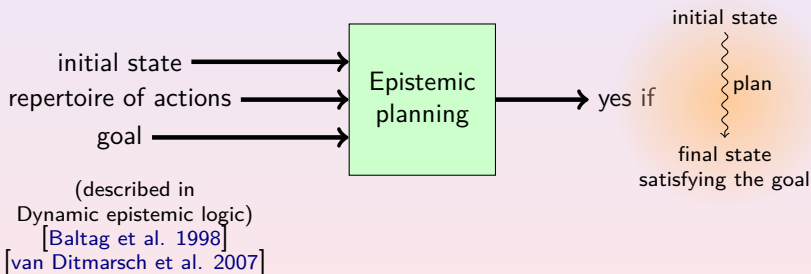
- Semi-product modal logics have high complexities;
[Gabbay et al. Many-Dimensional Modal Logics: Theory and Applications, 2003]
- Model checking more practical than theorem proving
[Halper, Vardi, KR 1991].

Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning**
 - Undecidability of epistemic planning
 - Decidability when pre/post are Boolean
 - Generalize to multi-player setting

Epistemic planning


[Andersen, Bolander, 2011]



Decidability and undecidability of epistemic planning

no postconditions:	✓	?	✗	✗	
Boolean postconditions:	✓	✗	✗	✗	
					→ modal depths of preconditions
	0	1	2	3	

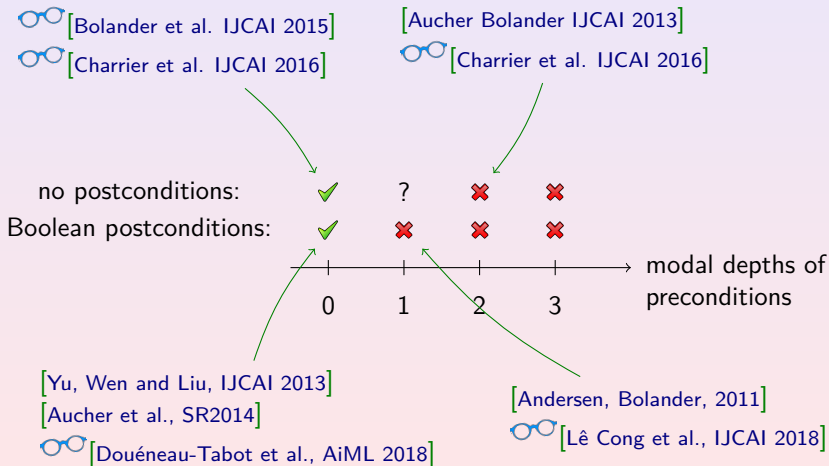
Decidability and undecidability of epistemic planning



Decidability and undecidability of epistemic planning

no postconditions:	✓	?	✗	✗
Boolean postconditions:	✓	✗	✗	✗
	0	1	2	3
	→ modal depths of preconditions			

Decidability and undecidability of epistemic planning



Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning**
 - **Undecidability of epistemic planning**
 - Decidability when pre/post are Boolean
 - Generalize to multi-player setting

Epistemic planning is undecidable

Theorem

Epistemic planning is undecidable for:

[Andersen, Bolander, JANCL 2011]

*two
agents*

+

*Boolean
post*

+

$md(pre) \leq 1$

+

*fixed
repertoire
of one action*

+

*6 atomic
propositions*




[Lê Cong, Pinchinat, __, IJCAI-ECAI 2018]

Proof: reduction from halting problem of a small universal cellular automaton.

Example: the 110 Rule cellular automaton




↓ time




Example: the 110 Rule cellular automaton



↓ time




Example: the 110 Rule cellular automaton




↓ time


Rules





Example: the 110 Rule cellular automaton




Rules




Example: the 110 Rule cellular automaton




Encoding an automaton configuration in a state




Encoding an automaton configuration in a state




Encoding an automaton configuration in a state




Encoding an automaton configuration in a state



Encoding an automaton configuration in a state




Encoding an automaton configuration in a state



Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 Epistemic planning**
 - Undecidability of epistemic planning
 - Decidability when pre/post are Boolean**
 - Generalize to multi-player setting

(Infinite) epistemic temporal structures



Epistemic planning: first-order query $\exists x, goal(x)$

Decidability when pre/post are Boolean

Theorem ([DBLP:conf/ijcai/YuWL13], [DBLP:journals/corr/AucherMP14])

When pre/post are Boolean, epistemic planning is *decidable*.

Epistemic planning is a first-order-query


first-order-query on **automatic** structures is **decidable**.

Epistemic temporal structures are **automatic**

Theorem ( [Douéneau-Tabot, Pinchinat and __, 2018])

Even decidable for goals in epistemic linear μ -calculus.

Automatic structure = defined by automatas



$$enc : \mathbb{N} \rightarrow \{1\}^*$$

$$n \mapsto 1^n$$


start \rightarrow  1

$\mathcal{A}_{\mathbb{N}}$



Example of an automatic structure

$\langle \mathbb{N}, \textit{iseven?}, \leq \rangle$




$\mathcal{A}_{\textit{iseven?}}$

Example of an automatic structure

$$S = \langle \mathbb{N}, \text{iseven?}, \leq \rangle$$


- $2 \leq 5$ iff "11 \leq 11111"
- $2 \leq 5$ iff word $\begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \begin{pmatrix} \square \\ 1 \end{pmatrix} \begin{pmatrix} \square \\ 1 \end{pmatrix} \begin{pmatrix} \square \\ 1 \end{pmatrix}$ is accepted by \mathcal{A}_{\leq}



Outline

- 1 Discussion about modeling actions
- 2 Formal definition of event models
- 3 Model checking
- 4 Theorem proving
- 5 **Epistemic planning**
 - Undecidability of epistemic planning
 - Decidability when pre/post are Boolean
 - **Generalize to multi-player setting**

Strategies



Definition

A *strategy* for player a is a function σ that maps any history $w e_1 \dots e_n$ to a deterministic epistemic action in the repertoire of a .

Definition

A *uniform strategy* for player a is a strategy σ such that

if $w e_1 \dots e_n \sim_a u e'_1 \dots e'_n$ then

$$\sigma(w e_1 \dots e_n) = \sigma(u e'_1 \dots e'_n)$$

Undecidability even for Boolean pre/post

Theorem

[Reif, Peterson, 1979] [Coulombe and Lynch, Def. 1, p. 14:7, FUN 2018]
[Maubert et al., IJCAI 2019]

The existence of uniform strategies for two players against an environment for achieving a goal φ is undecidable.

Decidability cases

public actions

[Belardinelli et al., 2017]

[Maubert et al., IJCAI 2019]



hierarchical information



[Maubert et al., 2018]

[Maubert et al., IJCAI 2019]



(picture idea from Raphael Berthon)

Complexity results on epistemic planning

	one centralized planner  [Bolander et al. IJCAI 2015]	many players  [Maubert et al., 2019]
public announcements	NP-c	PSPACE-c
public actions	PSPACE-c	EXPTIME-c
Boolean pre/post	decidable	undecidable [Reif, Peterson, 1979]
all	undecidable	

Uninformed semantics.


Perspectives: DEL and formal language theory

Question

Is epistemic planning one agent (pre md 1, ~~post~~) **decidable**?


First-order query
is decidable

Automatic structures



First-order query
is undecidable

Turing-complete
structures



Pushdown automata?
Causal hierarchy?

Perspectives

- Connection with logics for reasoning about strategies such as Alternating temporal-time logic, Strategy Logic, etc.

 [Maubert et al., 2019]

- Describing protocols/policies