

Logic, Automata, and Games in Linear Temporal Logics on Finite Traces

Giuseppe De Giacomo

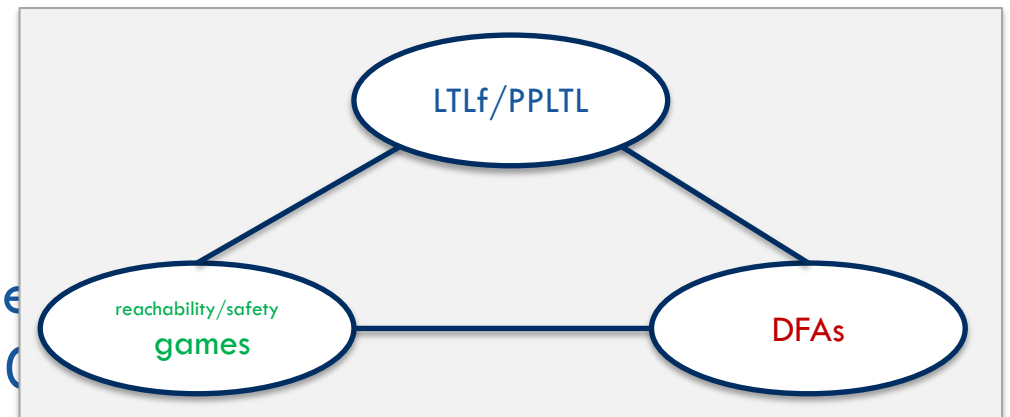
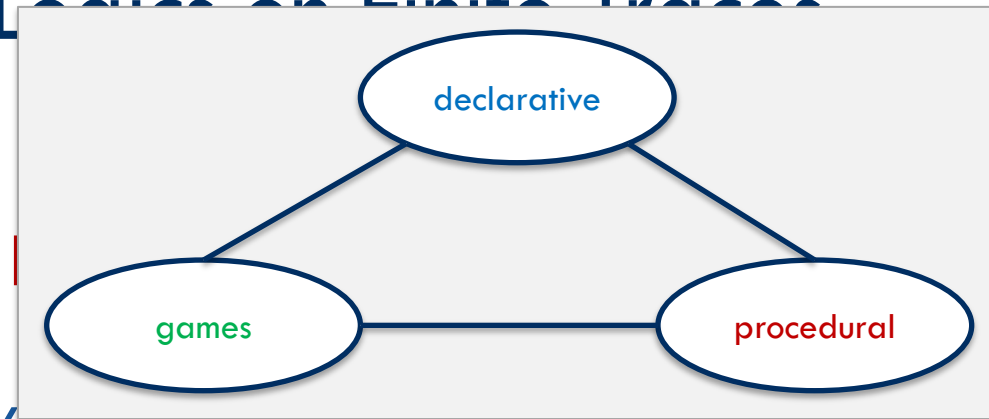
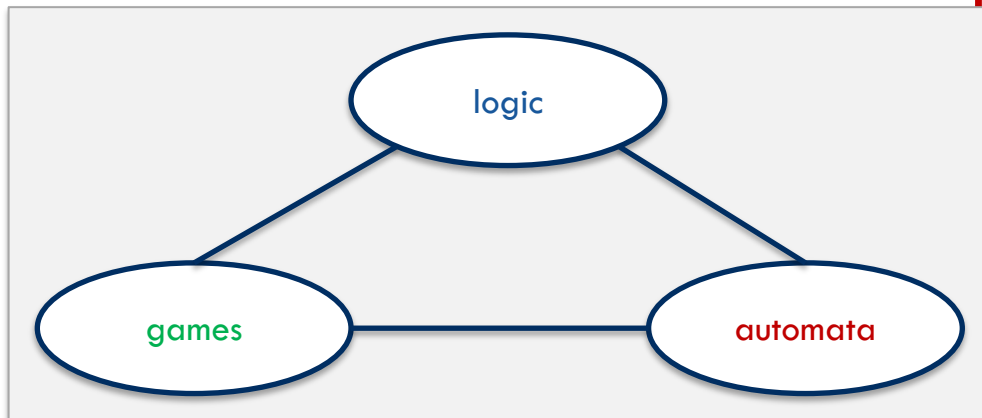
University of Oxford

GenPlan@NeurIPS2023 – New Orleans, USA
December 16, 2023

Logic, Automata, and Games in Linear Temporal Logic on Finite Traces

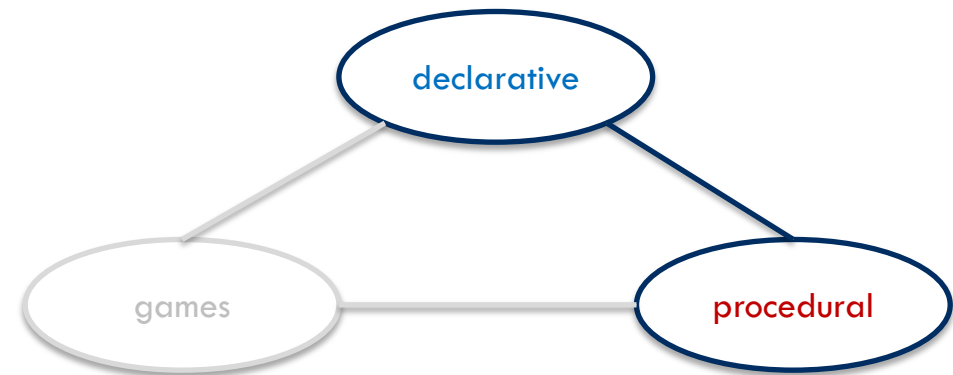
Giuseppe I...

University of Oxford

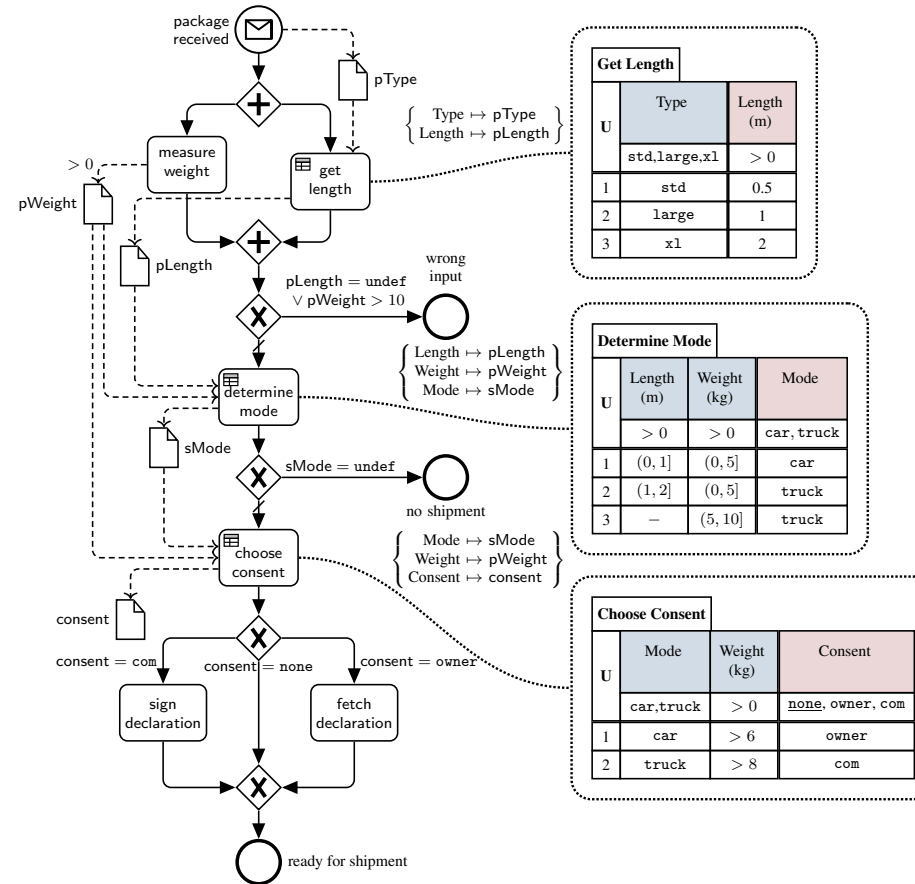


GenPlan@NeurIPS2023 – Ne
December 16, 20

DECLARATIVE TO PROCEDURAL

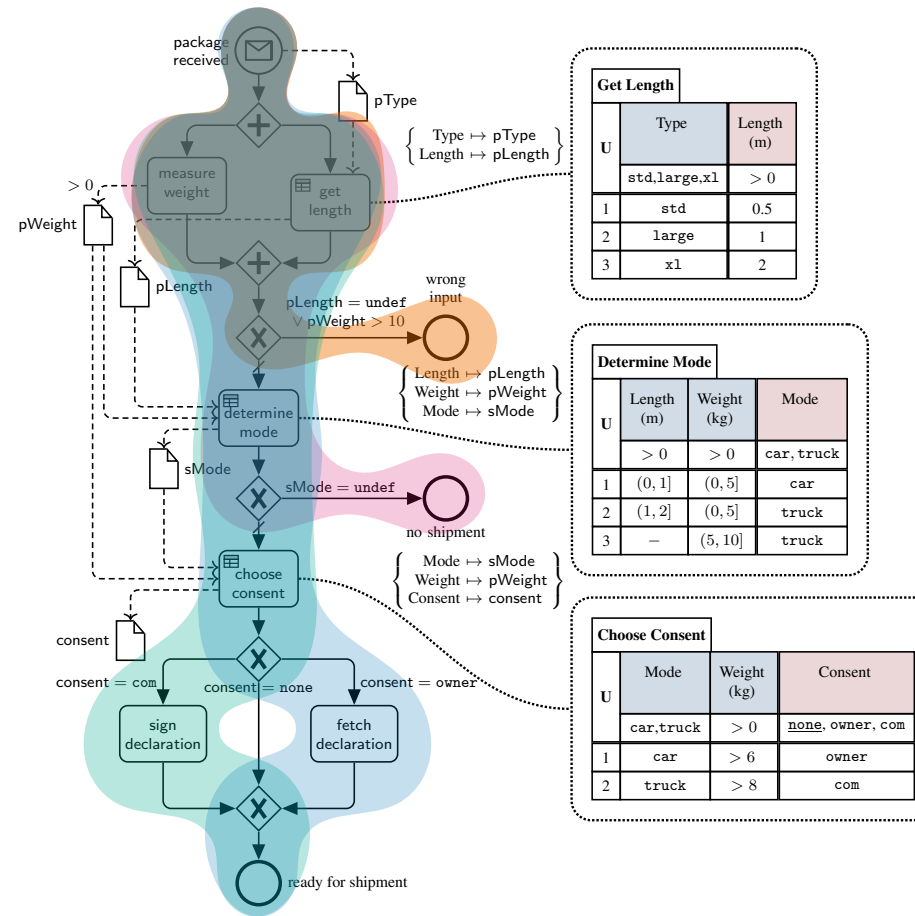


Procedural Specs



Courtesy of Marco Montali


Procedural Specs Induce Traces



Courtesy of Marco Montali

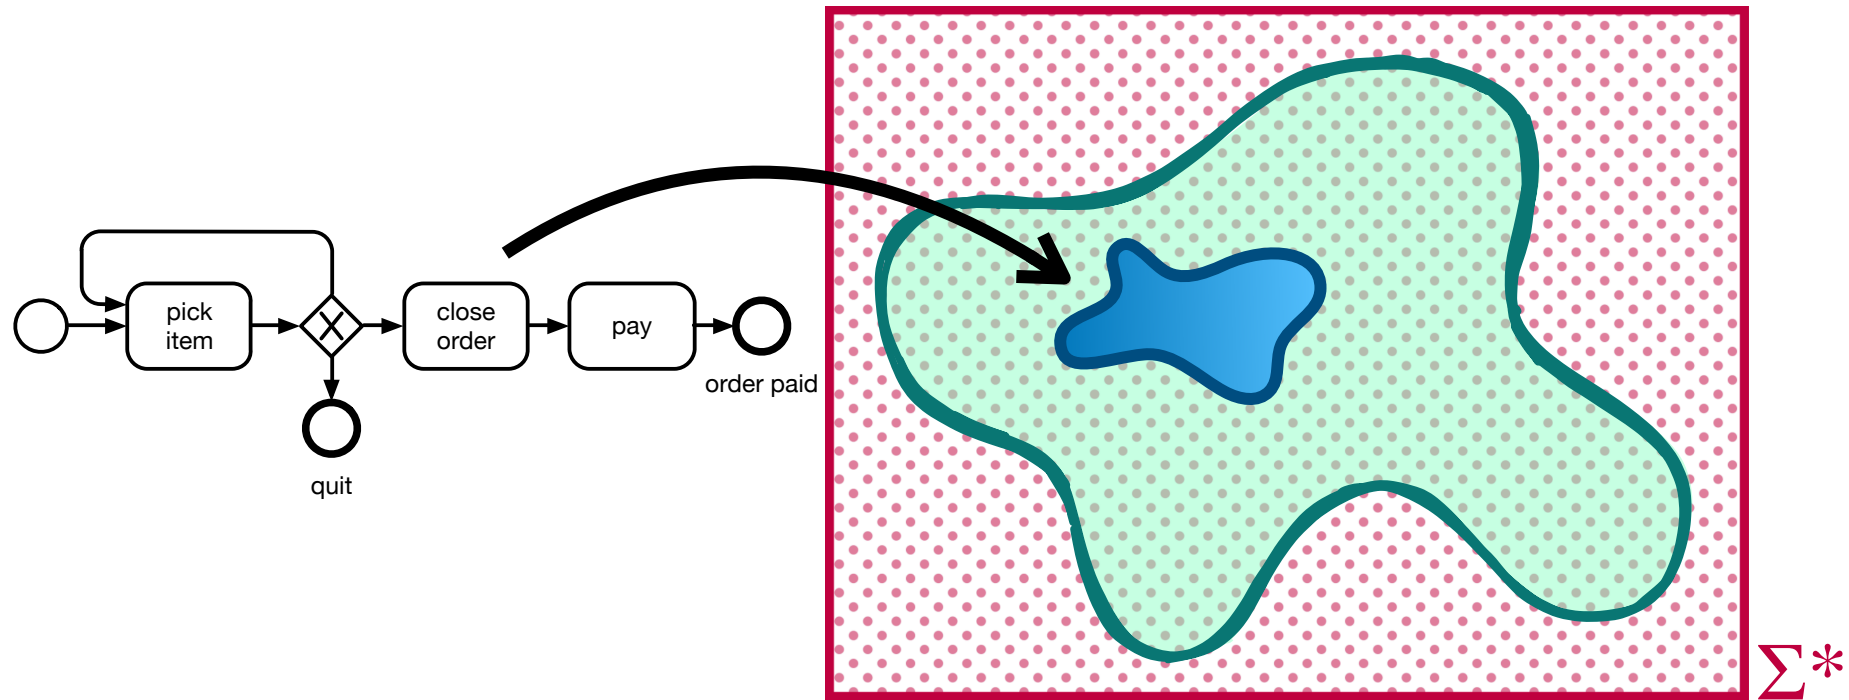
Traces Allowed by the Specs

A possibly infinite set of finite traces



Courtesy of Marco Montali

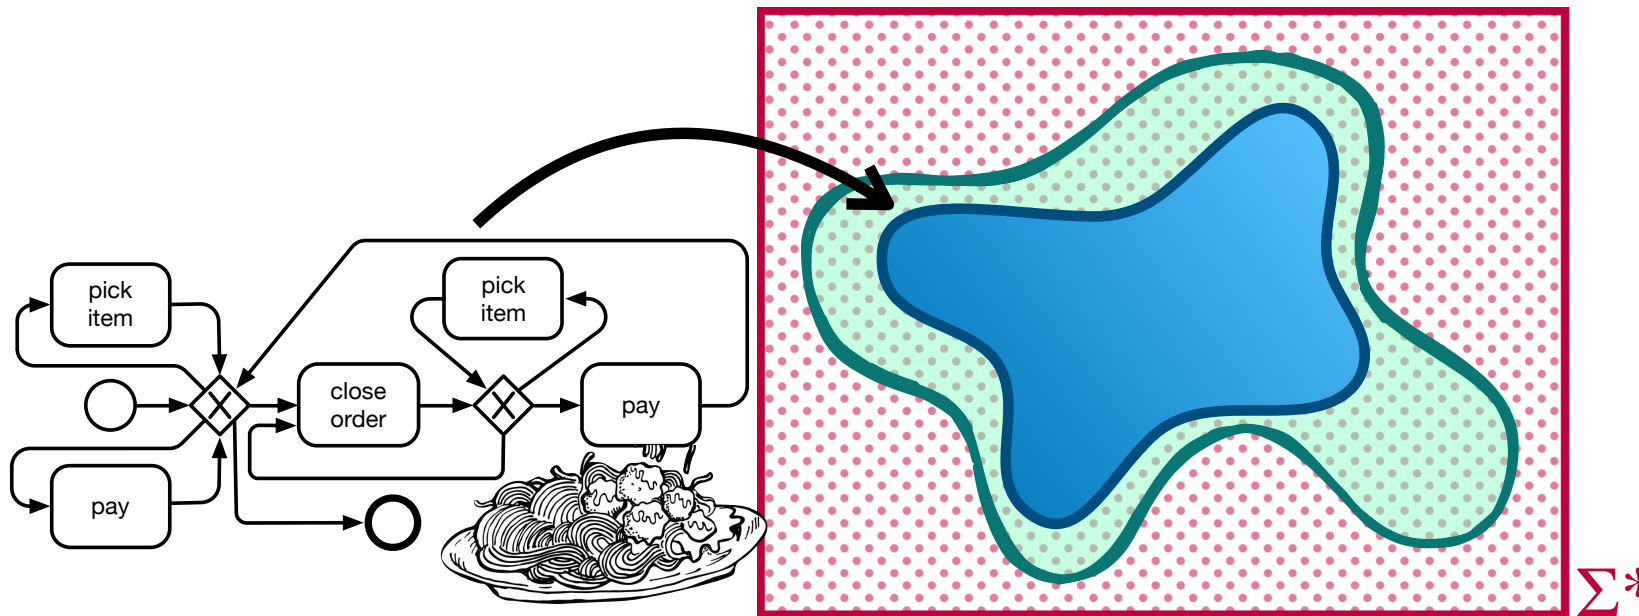
Traces Allowed by the Specs



Courtesy of Marco Montali


Traces Allowed by the Specs

Generalisation



Courtesy of Marco Montali

Reality Is Often More Flexible Than It Seems



Courtesy of Marco Montali

Declarative (Specifications of) Processes

Our goal



Compact
specification

represents



Reality

Courtesy of Marco Montali

Declarative Processes

In late 2000's the Business Process Management (BPM) community produced a brilliant idea:
 [PesicVanDerAalst06] [AlbertiEtAl06] [Montali2010] [PesicBovsnavkiVanDerAalst10]

The idea:


- Give the rules that a process should satisfy
- And nothing else!
- Extract the process from the rules only

Which specs?

- The one most used in formal methods for specifying process properties
- Linear Time Logic on Finite Traces (LTLf)

In other words:

- Drop explicit representation of process, and
- Use instead LTL formulas to specify the allowed process traces



Process

Imperative
model

Declarative
specification

Declarative Processes

In late 2000's the Business Process Management (BPM) community produced a brilliant idea:
 [PesicVanDerAalst06] [AlbertiEtAl06] [Montali2010] [PesicBovsnavkiVanDerAalst10]


The idea:

- Give the rules that a process should satisfy
- And nothing else!
- Extract the process from the rules only

In other words: Automatically synthesize process from declarative specs

It can be seen as the fulfillment of the CS dream:

- Devise a technique for the “mechanical translation of human-understandable task specifications to a program that is known to meet the specifications.” [Vardi - The Siren Song of Temporal Synthesis 2018] (more later)



Process

Imperative
model

Declarative
specification

Declarative Processes

Originally a controlled set of notable LTL formulas on were proposed for process specification (and a suitable graphical notation provided) [PesciVanDerAalst06]

Can we use any LTL formula as a declarative spec?

Example (Main DECLARE Patterns)

NAME	NOTATION	LTL _f	DESCRIPTION
Existence		$\diamond a$	a must be executed at least once
Resp. existence		$\diamond a \supset \diamond b$	If a is executed, then b must be executed as well
Response		$\square(a \supset \diamond b)$	Every time a is executed, b must be executed afterwards
Precedence		$\neg b \mathcal{W} a$	b can be executed only if a has been executed before
Alt. Response		$\square(a \supset \bigcirc(\neg a \cup b))$	Every a must be followed by b, without any other a in in
Chain Response		$\square(a \supset \bigcirc b)$	If a is executed then b must be executed next
Chain Precedence		$\square(\bigcirc b \supset a)$	Task b can be executed only immediately after a
Not Coexistence		$\neg(\diamond a \wedge \diamond b)$	Only one among tasks a and b can be executed
Neg. Succession		$\square(a \supset \neg \diamond b)$	Task a cannot be followed by b, and b cannot be preceded by a
Neg. Chain Succ.		$\square(a \supset \bigcirc \neg b)$	Tasks a and b cannot be executed next to each other

No! What if the spec is: *always eventually Happy* ?

Yes, if you focus on finite trace! (Specs in LTL_f instead of LTL)


Assumes only one activity (proposition) true at each

G. De Giacomo, R. De Masellis, M. Montali: Reasoning on LTL on Finite Traces: Insensitivity to Infiniteness. AAAI 2014

FOCUS ON FINITE TRACES

Formal Methods

- Rigorous guarantees about the behavior of computational systems
- Wide-spread industrial adoption
- Main tasks
 - **Formalisms for specs of dynamic properties:**
 - E.g., Linear Temporal Logic
 - **Verification:**
 - Check if the system satisfies specs.
 - **Synthesis:**
 - Synthesize a system that satisfies specs.



AIRBUS
GROUP






aws **AUTOMATED**
REASONING GROUP

Model Checking vs Synthesis

The “big bang” of the application of temporal logic to program verification: **Linear Temporal Logic (LTL)** (Pnueli, 1977)

Semantics of LTL is over ω -words, i.e., infinite traces

Note:

- ω -automata algorithms scale well as long as you do not have to determinize 
- For synthesizing strategies/policy determinization is essential 
- In AI are more interested in finite-trace semantics 

Focus on Finite Traces is Shared by AI

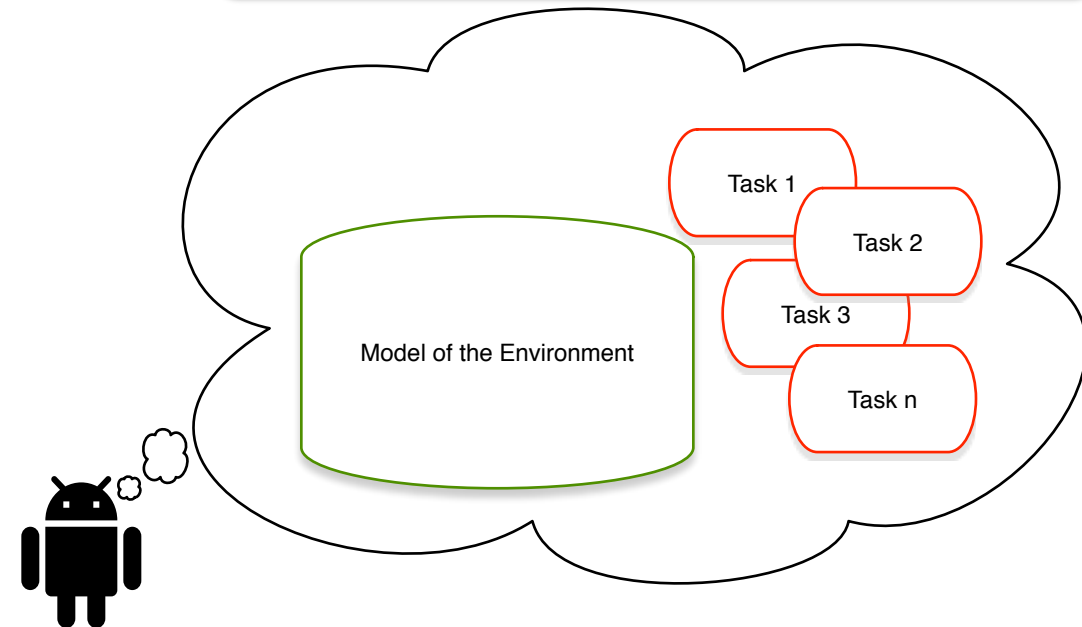
Planning in AI:

- Is all about having a task specification or “goal” and producing a “plan” (or strategy or policy) to satisfy the task in the environment model.
- Which tasks?
 - A **task that terminates!**
 - Typically, just reaching a certain state in the environment

Why tasks that terminates?

- Because it is the agent that is planning/reasoning
- If the task would not terminate, the agent would be stuck into doing the same task forever
- But then, why bother with equipping it with a model of the environment and of the task at all?
- Note it is the agent, NOT the designer, who has such a model

- *In Formal Methods focus on infinite traces*
- *In AI focus on finite traces LTL \rightarrow LTLf*



Linear Time Temporal Logics on Finite Traces

LTL_f/LDL_f : linear temporal logics on finite traces [DeGiacomoVardi2013]

LTL_f : linear time temporal logic on finite traces

Same syntax as standard LTL but interpreted over finite traces

$$\varphi ::= A \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \text{next } \varphi \mid \text{eventually } \varphi \mid \text{always } \varphi \mid \varphi_1 \text{ until } \varphi_2$$

Examples:	<i>eventually A</i>	“eventually A”	<i>reachability</i>
	<i>always A</i>	“always A”	<i>safety</i>
	<i>always(A → eventually B)</i>	“always if A then eventually B”	<i>reactiveness</i>
	<i>A until B</i>	“A until B”	<i>until</i>
	<i>¬B until A ∨ always ¬B</i>	“A before B”	<i>precedence</i>

LDL_f : linear dynamic logic on finite traces

Same syntax as PDL but interpreted over finite traces

$$\varphi ::= tt \mid A \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \langle \rho \rangle \varphi \mid [\rho] \varphi \quad \rho ::= A \mid \varphi? \mid \rho_1 + \rho_2 \mid \rho_1; \rho_2 \mid \rho^*$$

Adds the possibility of expressing procedural constraints/goals [Reiter01], [BaierFritzMcIlraith07]:

$$\delta ::= A \mid \varphi? \mid \delta_1 + \delta_2 \mid \delta_1; \delta_2 \mid \delta^* \mid \text{if } \phi \text{ then } \delta_1 \text{ else } \delta_2 \mid \text{while } \phi \text{ do } \delta$$

where **if** and **while** are abbreviations: **if** ϕ **then** δ_1 **else** $\delta_2 \doteq (\phi?; \delta_1) + (\neg\phi?; \delta_2)$ and **while** ϕ **do** $\delta \doteq (\phi?; \delta)^*; \neg\phi?$

Linear Time Temporal Logics on Finite Traces

Example

- “All coffee requests from person p will eventually be served”:

$$\text{always}(\text{request}_p \rightarrow \text{eventually coffee}_p) \quad [\text{true}^*](\text{request}_p \rightarrow \langle \text{true}^* \rangle \text{coffee}_p)$$

- “Every time the robot opens door d it closes it immediately after”:

$$\text{always}(\text{openDoor}_d \rightarrow \text{next closeDoor}_d) \quad [\text{true}^*](\langle \text{openDoor}_d \rangle \text{closeDoor}_d)$$

- “Before entering restricted area a the robot must have permission for a ”:

$$\neg \text{inArea}_a \text{ until } \text{getPerm}_a \vee \text{always } \neg \text{inArea}_a \quad \langle (\neg \text{inArea}_a)^* \rangle \text{getPerm}_a \vee [\text{true}^*] \neg \text{inArea}_a$$

- “Each time the robot enters the restricted area a it must have a new permission for a ”:

$$\langle (\neg \text{inArea}_a^* ; \text{getPerm}_a ; \neg \text{inArea}_a^* ; \text{inArea}_a ; \text{inArea}_a^*)^* ; \neg \text{inArea}_a^* \rangle \text{end}$$

- “At every point, if it is hot then, if the air-conditioning system is off, turn it on, else don’t turn it off”:

$$[\text{true}^*] \langle \text{if } (\text{hot}) \text{ then} \\ \text{if } (\neg \text{airOn}) \text{ then } \text{turnOnAir} \\ \text{else } \neg \text{turnOffAir} \rangle \text{true}$$

LTLf to Automata

DFA are indeed machines and hence processes!

Key point

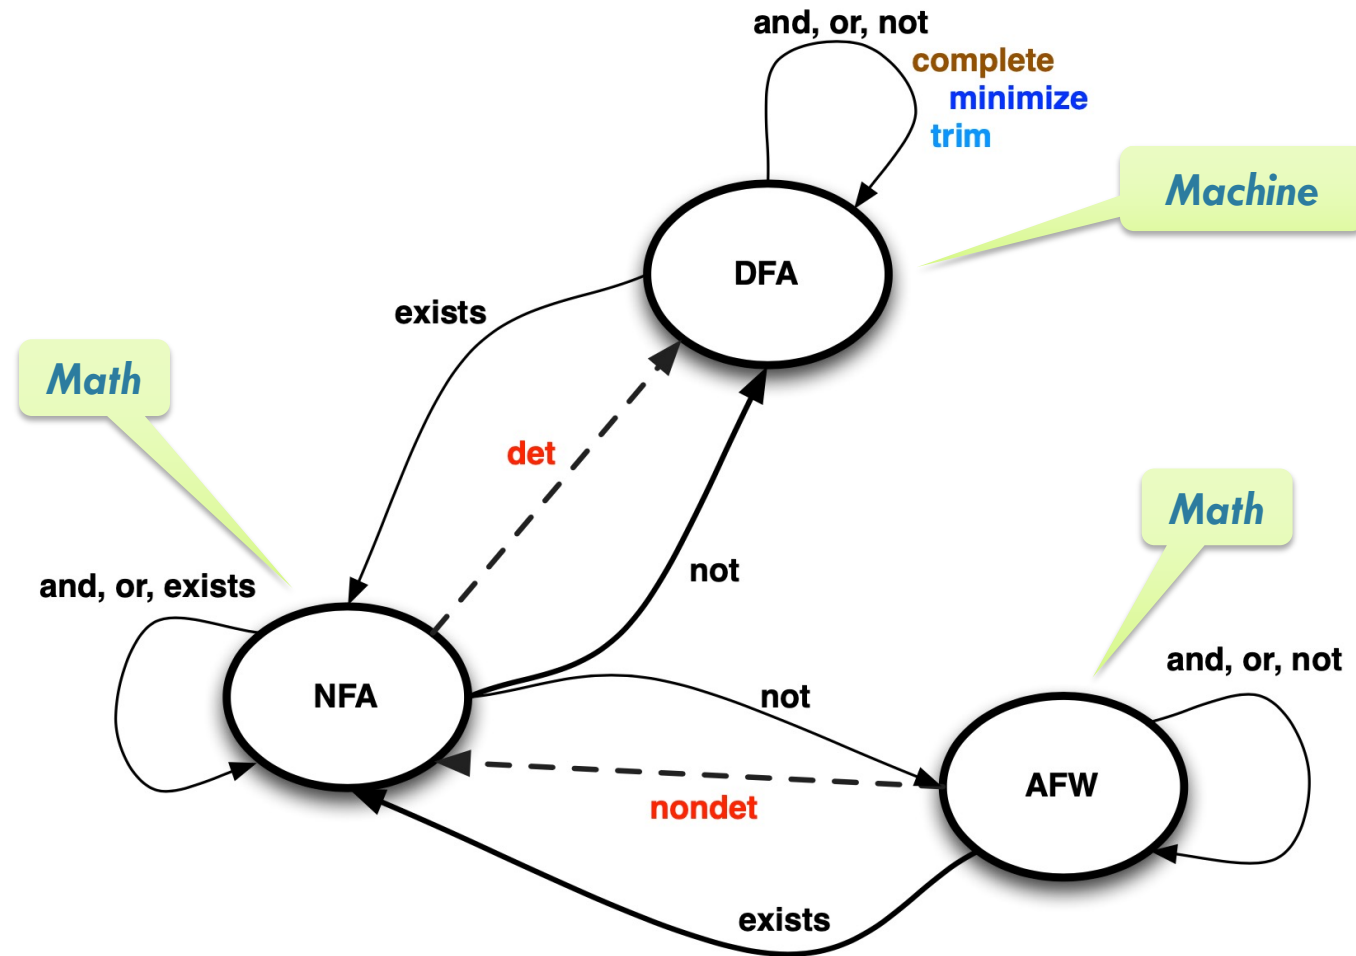
LTL_f/LDL_f formulas can be translated into a finite-state automaton on finite words \mathcal{A}_φ such that:

$$t \models \varphi \text{ iff } t \in \mathcal{L}(\mathcal{A}_\varphi)$$

- in **linear time** if \mathcal{A}_φ is an **Alternating Automata (AFW)**;
- in **exponential time** if \mathcal{A}_φ is an **Nondeterministic Finite-state Automaton (NFA)**;
- in **double exponential time** if \mathcal{A}_φ is an **Deterministic Finite-state Automaton (DFA)**.

We can compile reasoning into automata based procedures!

Regular Automata



LTLf to Automata

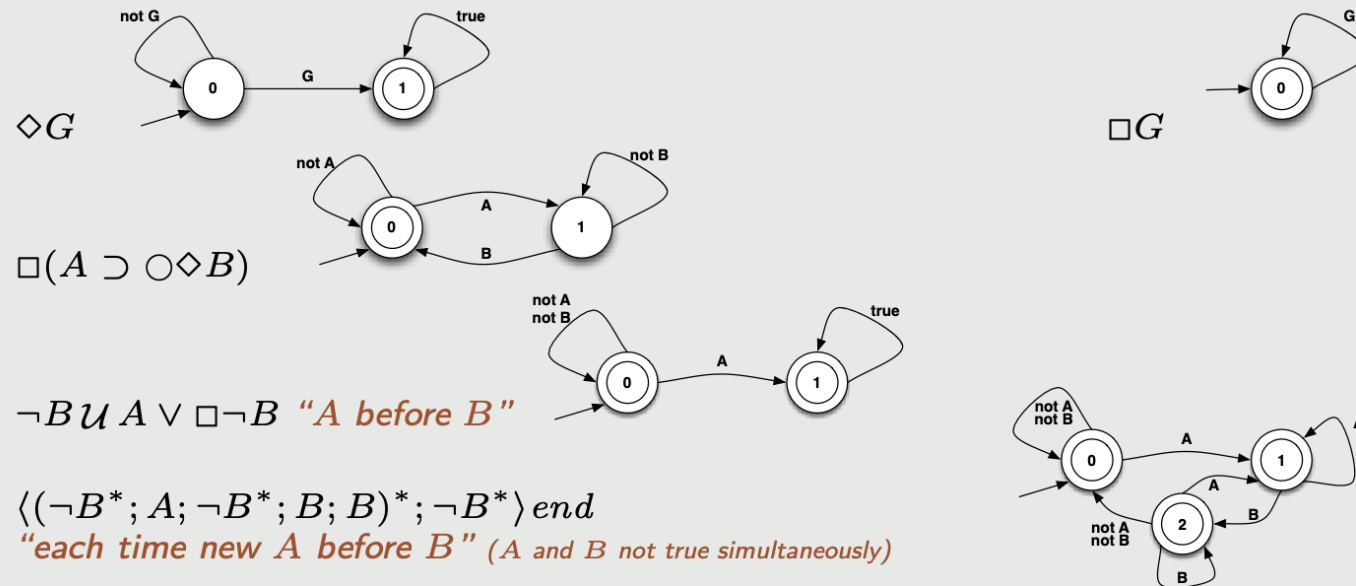
Key point

LTL_f/LDL_f formulas can be translated into deterministic finite state automata (DFA).

$$t \models \varphi \text{ iff } t \in \mathcal{L}(A_\varphi)$$

where A_φ is the DFA φ is translated into.

Example (Automata for some LTL_f/LDL_f formulas)



(online software for LTL_f2DFA : <http://ltlf2dfa.diag.uniroma1.it>)
 (online software for LDL_f2DFA : <http://lydia.whitemech.it>)

Pure Past LTL

PPLTL is a variant of LTLf that looks at traces backward (from now toward the past)

PPLTL: Pure Past LTL

$$\varphi ::= A \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \Theta\varphi \mid \varphi_1 \mathcal{S} \varphi_2$$

- PPLTL has the same expressive power of LTLf.
- The DFA \mathcal{A}_Φ of a PPLTL formula Φ is worst-case single exponential (vs double-exponential for LTLf):
 - ▶ build AFA reading the trace backward (same AFA as for LTLf - linear)
 - ▶ compute the DFA of the reverse language (single exponential)
- Best algorithm for systematic translation between LTLf and PPLTL is 3EXPTIME,...
- ... but in several significant cases is polynomial, e.g., all DECLARE patterns [GeattiMontaliRivkinArXiv2022].

Pure Past LTL

Actually PPLTL formulas can be evaluated in a dynamic programming method using only two instants:

- Now
- One step in the past (prev)

Current and Previous Instants

PPLTL formulas can be evaluated considering only the current and previous instants:

$$\begin{aligned}
 eval(A, \Pi_{now}, \Pi_{prev}) &= A \text{ if variable "A" is true in } \Pi_{now} \\
 eval(\Theta\varphi, \Pi_{now}, \Pi_{prev}) &= \Theta\varphi \text{ if variable "\Theta\varphi" is true in } \Pi_{prev} \\
 eval(\neg\varphi, \Pi_{now}, \Pi_{prev}) &= \neg eval(\varphi, \Pi_{now}, \Pi_{prev}) \\
 eval(\varphi_1 \wedge \varphi_2, \Pi_{now}, \Pi_{prev}) &= eval(\varphi_1, \Pi_{now}, \Pi_{prev}) \wedge eval(\varphi_2, \Pi_{now}, \Pi_{prev}) \\
 eval(\varphi_1 \vee \varphi_2, \Pi_{now}, \Pi_{prev}) &= eval(\varphi_1, \Pi_{now}, \Pi_{prev}) \vee eval(\varphi_2, \Pi_{now}, \Pi_{prev}) \\
 eval(\varphi_1 \mathcal{S} \varphi_2, \Pi_{now}, \Pi_{prev}) &= eval([\varphi_2 \vee (\varphi_1 \wedge \Theta(\varphi_1 \mathcal{S} \varphi_2))], \Pi_{now}, \Pi_{prev})
 \end{aligned}$$

Remember that the fixpoint equation for $\varphi_1 \mathcal{S} \varphi_2$ is
 $\varphi_1 \mathcal{S} \varphi_2 \equiv \varphi_2 \vee (\varphi_1 \wedge \Theta(\varphi_1 \mathcal{S} \varphi_2)).$

Pure Past LTL

As a result one can built a symbolic DFA in linear time!

Symbolic DFA

$$\mathcal{A}_\Phi = (AP, PP, \Pi_{prev}^0, Trans, Final)$$

where

- AP set of boolean variable, one for each atomic proposition, representing which propositions are true/false in the current instant;
- PP set of boolean variable, one component for each sub-formulas the form $\ominus\varphi$ in Fisher-Ladner closure of Φ , representing which formulas were true/false at the previous instant;
- $\Pi_{prev}^0 = (false, \dots, false)$, initially all formulas of the form $\ominus\varphi$ are false
- $Trans(\Pi_{now}, \Pi_{prev}) = Trans_i(\Pi_{now}, \Pi_{prev}) \times \dots \times Trans_n(\Pi_{now}, \Pi_{prev})$ with $n = |\Pi_{prev}|$, lat where for each variable " $\ominus\varphi_i$ ":

$$Trans_i(\Pi_{now}, \Pi_{prev}) = eval(\varphi_i, \Pi_{now}, \Pi_{prev})$$

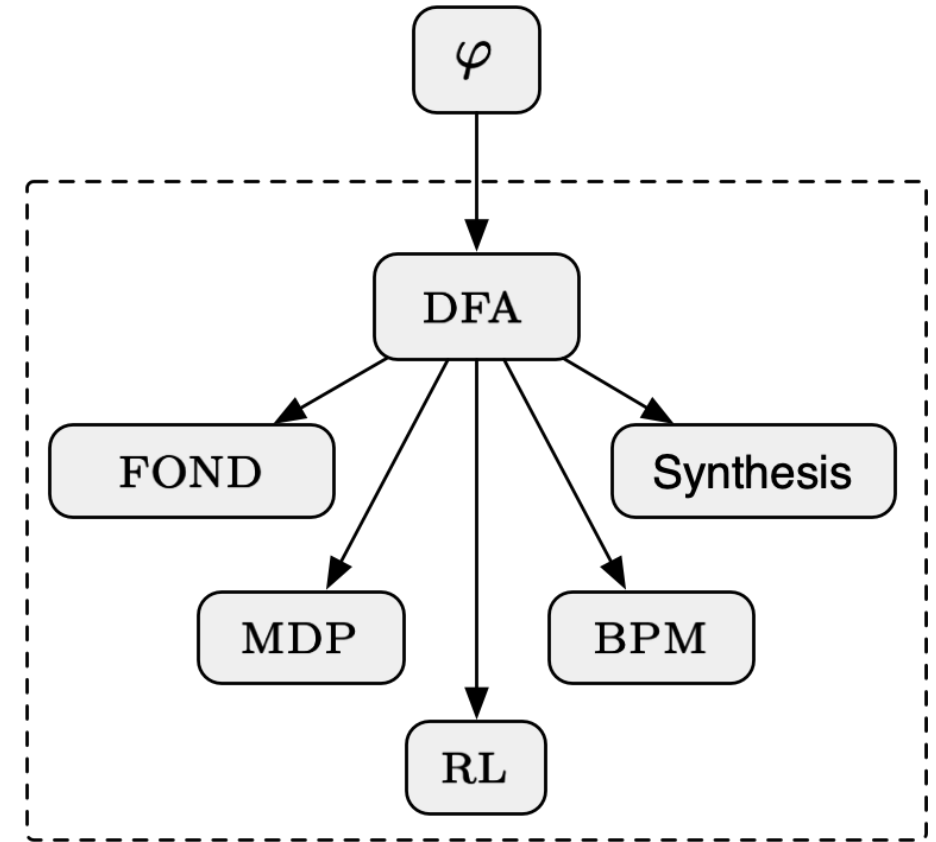
- $Final(\Pi_{now}, \Pi_{prev}) \equiv eval(\Phi, \Pi_{now}, \Pi_{prev})$

Important: \mathcal{A}_Φ is linear in Φ .


Several Applications of LTLf/PPLTL Specs

Many Applications:

- Planning for temporally extended goals
- Several forms of Synthesis
- MDP with non-Markovian rewards
- Reinforcement Learning for non-Markovian tasks
- Declarative Process Specification in BPM



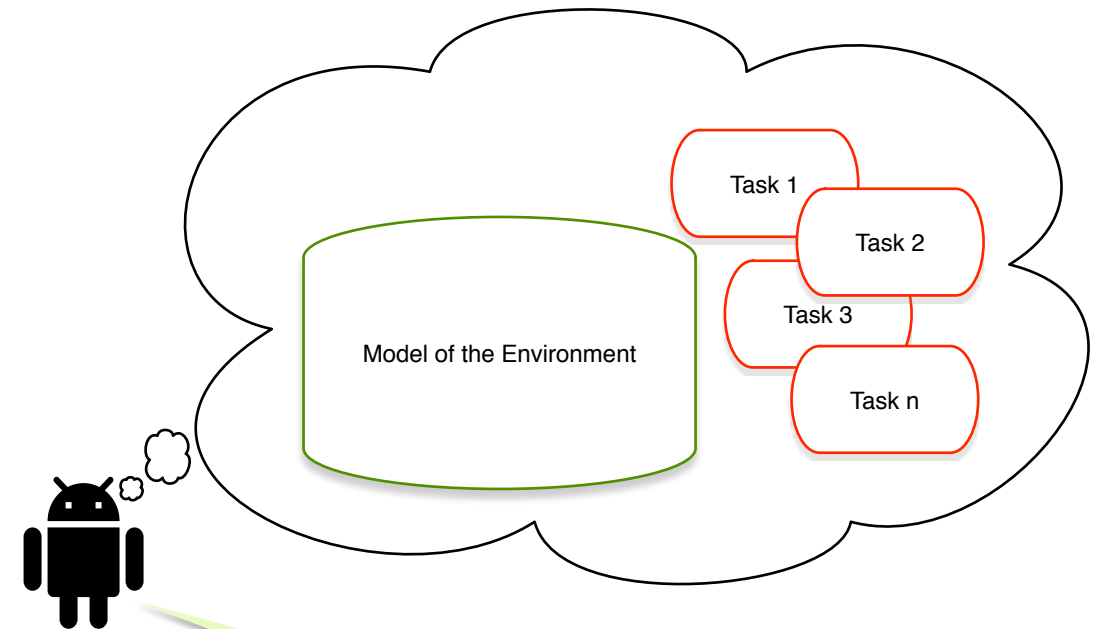
DECLARATIVE TO PROCEDURAL TO GAMES



Specify Models and Tasks as Processes in Formal Methods

- **Environment Model (ENV)**
 - Spec. of **environment** possible behaviors (**ENV**)
 - Think of each **behavior** as a choice function resolving nondeterminism of a nondeterministic domain
- **ENV** expressed as
 - **nondeterministic planning domains**
 - **LTL/LTLf** specifications
- **Agent's Task (TASK/GOAL)**
 - Spec. of **agent's task**
 - **TASK** expressed in **LTL/LTLf**
- **Find agent's plan/behavior/policy/strategy** that fulfills **TASK** against **all** behaviors of **ENV**

Specify **ENV** and **TASK** as with formalism used in **Formal Methods** e.g., **LTL/LTLf**



Find agent's **strategy** that is **winning against ENV**

Nondeterministic Domains as Deterministic Automata

Nondeterministic domain (including initial state)

$\mathcal{D} = (2^{\mathcal{F}}, \mathcal{A}, s_0, \delta, \alpha)$ where:

- \mathcal{F} **fluents** (atomic propositions)
- \mathcal{A} **actions** (atomic symbols)
- $2^{\mathcal{F}}$ set of states
- s_0 initial state (initial assignment to fluents)
- $\alpha(s) \subseteq \mathcal{A}$ represents **action preconditions**
- $\delta(s, a, s')$ with $a \in \alpha(s)$ represents **action effects (including frame)**.

Automaton $A_{\mathcal{D}}$ for \mathcal{D} is a **DFA!!!**

$A_{\mathcal{D}} = (2^{\mathcal{F} \cup \mathcal{A}}, (2^{\mathcal{F}} \cup \{s_{init}\}), s_{init}, \rho, F)$ where:

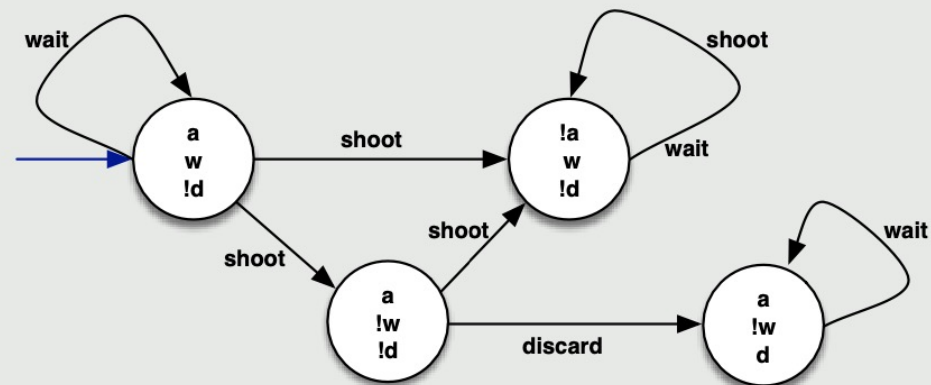
- $2^{\mathcal{F} \cup \mathcal{A}}$ alphabet (actions \mathcal{A} include dummy *start* action)
- $2^{\mathcal{F}} \cup \{s_{init}\}$ set of states
- s_{init} dummy initial state
- $F = 2^{\mathcal{F}}$ (all states of the domain are final)
- $\rho(s, [a, s']) = s'$ **with** $a \in \alpha(s)$, **and** $\delta(s, a, s')$ $\rho(s_{init}, [start, s_0]) = s_0$

(notation: $[a, s']$ stands for $\{a\} \cup s'$)

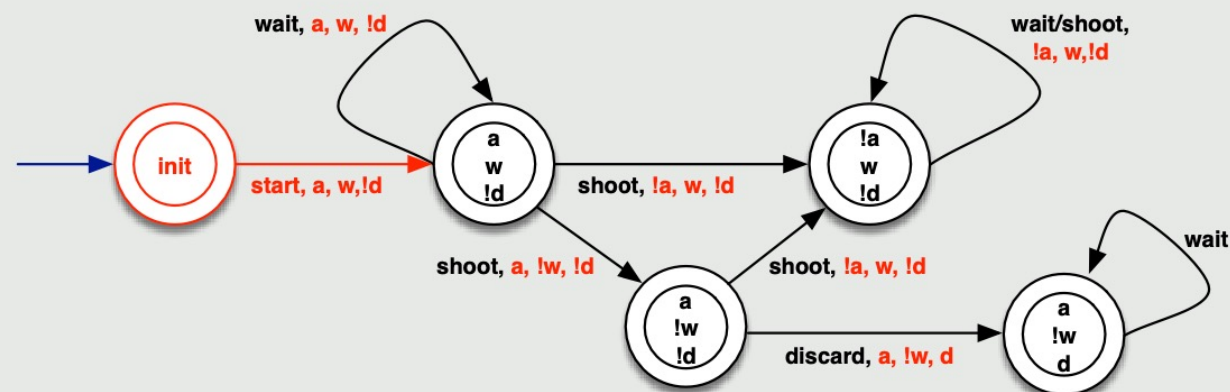
Nondeterministic Domains as Deterministic Automata

Example (Simplified Yale shooting domain)

- Domain \mathcal{D} :



- DFA $A_{\mathcal{D}}$:

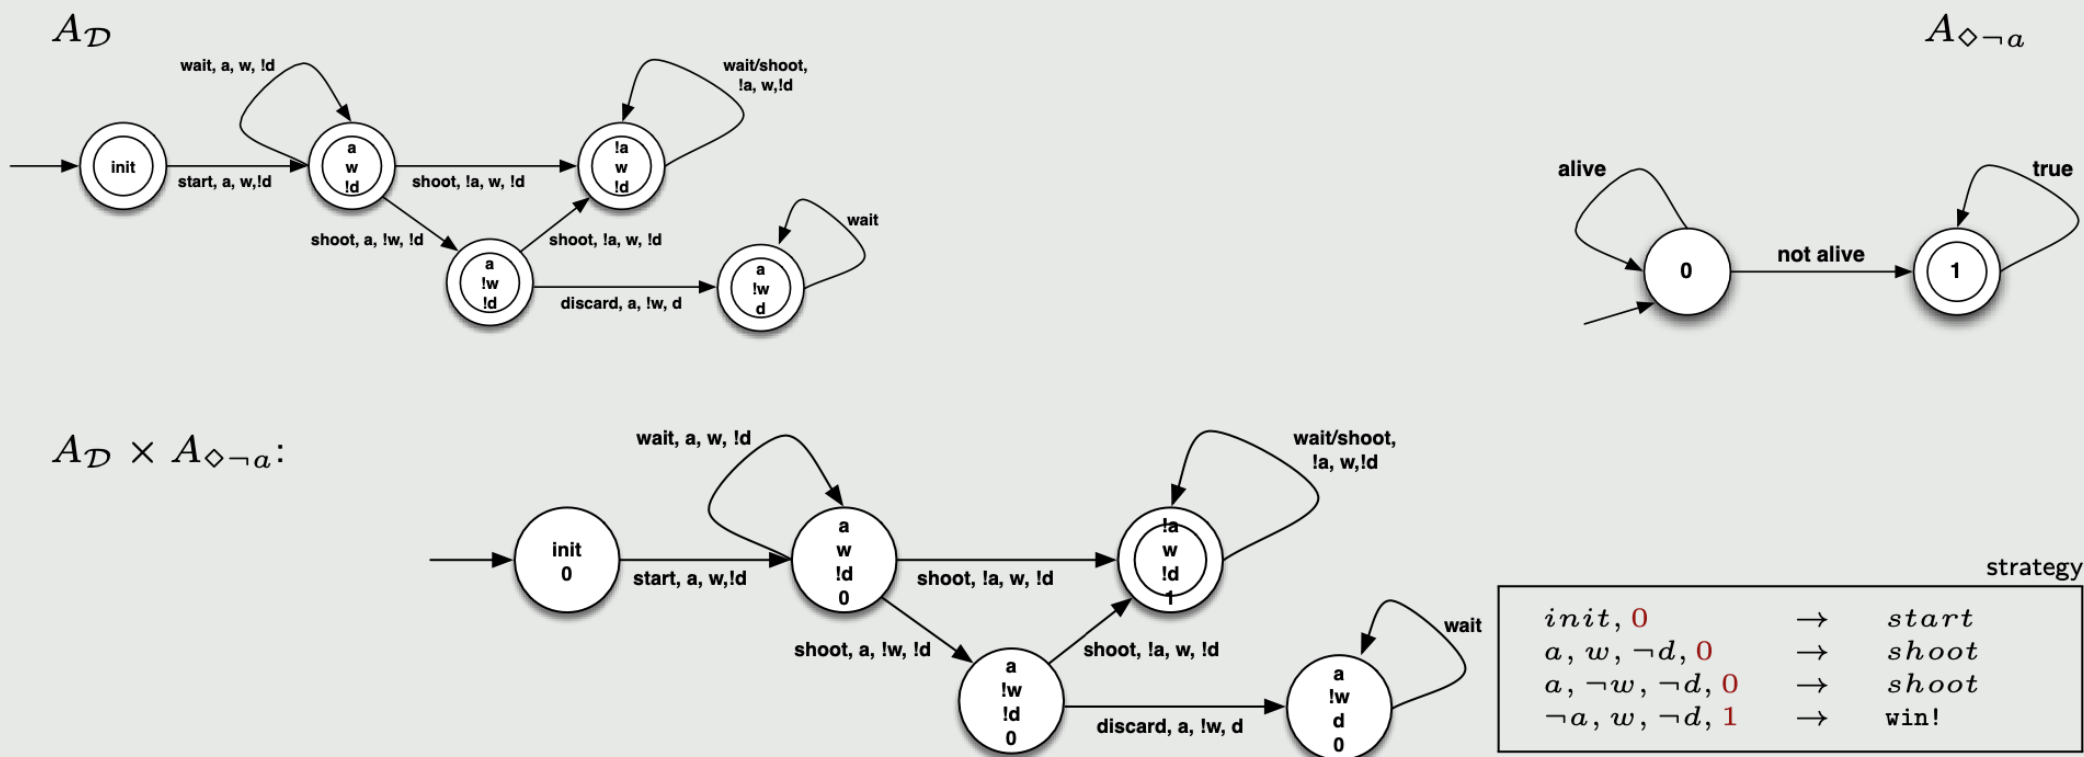


Planning in Nondeterministic Domains for Reachability Goals

Planning in nondeterministic domains

- Set the **arena** formed by all traces that satisfy both the DFA $A_{\mathcal{D}}$ for \mathcal{D} and the DFA for $\diamond G$ where G is the goal.
- Compute a **winning strategy**. *(EXPTIME-complete in \mathcal{D} , constant in G)*


Example (Simplified Yale shooting domain)




Planning in Nondeterministic Domains for Arbitrary LTLf Goals

In general, we need first to determinize the NFA for LTL_f/LDL_f formula

NFA for $\diamond \square \neg a$


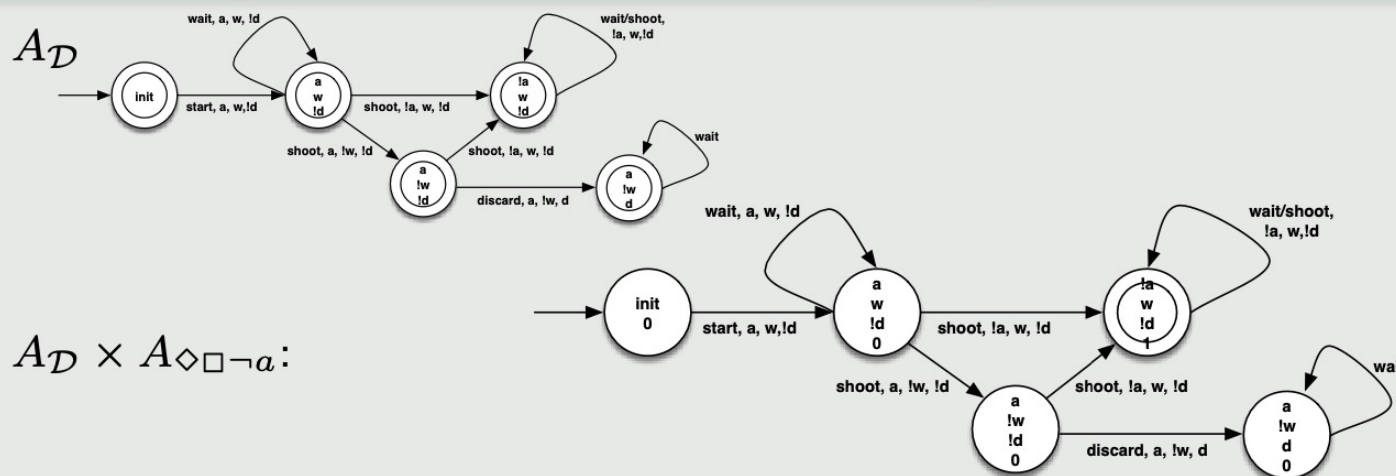


corresponding DFA



(DFA can be exponential in NFA in general)

Example (Simplified Yale shooting domain)



strategy

$init, 0$	\rightarrow	$start$
$a, w, \neg d, 0$	\rightarrow	$shoot$
$a, \neg w, \neg d, 0$	\rightarrow	$shoot$
$\neg a, w, \neg d, 1$	\rightarrow	$win!$

Planning in Nondeterministic Domains for LTLf Goals

DFA games

A **DFA game** $\mathcal{G} = (2^{\mathcal{F} \cup \mathcal{A}}, S, s_{init}, \varrho, F)$, is such that:

- \mathcal{F} controlled by **environment**; \mathcal{A} controlled by **agent**;
- $2^{\mathcal{F} \cup \mathcal{A}}$, alphabet of game;
- S , states of game;
- s_{init} , initial state of game;
- $\varrho : S \times 2^{\mathcal{F} \cup \mathcal{A}} \rightarrow S$, transition function of the game: given current state s and a choice of action a and resulting fluents values E the resulting state of game is $\varrho(s, [a, E]) = s'$;
- F , final states of game, where game can be considered terminated.

*This is the **game area**, in which agent and env will play!*

*It is **not only the domain**, but it is obtained from the **domain** and the **DFA of the formula***

Winning Strategy:

- A play is **winning** for the agent if such a play leads from the initial to a final state.
- A **strategy** for the agent is a function $f : (2^{\mathcal{F}})^* \rightarrow \mathcal{A}$ that, given a **history of choices from the environment**, decides which action \mathcal{A} to do next.
- A **winning strategy** is a strategy $f : (2^{\mathcal{F}})^* \rightarrow \mathcal{A}$ such that for all traces π with $a_i = f(\pi_{\mathcal{F}}|_i)$ we have that π leads to a final state of \mathcal{G} .

Planning in Nondeterministic Domains for LTLf Goals

Winning states for DFA games

Let denote the **set of final states** of \mathcal{G} as:

$$\llbracket F \rrbracket = \{s \in \mathcal{S} \mid s \models F\}$$

and let's define the **(adversarial preimage of a set \mathcal{E})** the following function:

$$PreAdv(\mathcal{E}) = \{s \in \mathcal{S} \mid \exists a \in \alpha(s). \forall s' \in \mathcal{S}. \delta(s, a, s') \supset s' \in \mathcal{E}\}$$

Compute the set Win of winning states of DFA game, i.e., states from which the agent can reach the final states F , by **least-fixpoint**:

- $Win_0 = \llbracket F \rrbracket$ (the final states)
- $Win_{i+1} = Win_i \cup PreAdv(Win_i)$
- $Win = \bigcup_i Win_i$

```

Wold := ∅
W := ⌊F⌋
while (W ≠ Wold) {
    Wold := W
    W := W ∪ PreAdv(W)
}
return W
    
```

(Computing Win is linear in the number of states in \mathcal{G})

Computing the winning strategy

Let's define $\omega : \mathcal{S} \rightarrow 2^{\mathcal{A}}$ as: $\omega(s) = \{a \in \alpha(s) \mid \text{if } s \in Win_{i+1} - Win_i \text{ then } \forall s'. \delta(s, a, s') \supset s' \in Win_i\}$

- **Every way** of restricting $\omega(s)$ to return only one action (chosen arbitrarily) gives a **winning strategy** for \mathcal{G} .
- Note s is a **state of the game!** not of the domain only!
To phrase ω wrt the domain only, we need to return a **stateful transducer** with transitions from the game.

Planning in Nondeterministic Domains for LTLf Goals

FOND for LTL_f goals

Algorithm: FOND for LTL_f/LDL_f goals

- 1: Given a FOND domain \mathcal{D} and an LTL_f/LDL_f goal φ
- 2: Compute DFA A_φ for φ (double exponential)
- 3: Compute product of \mathcal{D} and A_φ (polynomial)
- 4: Synthesize winning strategy for DFA game (linear)
- 5: Return strategy

It's a game agent vs env!

- Build the arena
- Play to win!

Theorem ([DeGiacomoRubinIJCAI18])

FOND for LTL_f/LDL_f goals is:

- EXPTIME-complete in the domain (assuming a logarithmic representation as in PDDL);
- 2EXPTIME-complete in the goal.

Same as classic FOND

Note we have separated costs in the model (DOM) and the task GOAL!
(c.f. data vs query complexity in Databases)

Planning in Nondeterministic Domains for PPLTL Goals

FOND for LTL_f goals

Algorithm: FOND for LTL_f , PPLTL goals

- 1: Given a FOND domain \mathcal{D} and an LTL_f/LDL_f goal φ
- 2: Compute DFA A_φ for φ (double exponential) *single exp for PPLTL*
- 3: Compute product of \mathcal{D} and A_φ (polynomial)
- 4: Synthesize winning strategy for DFA game (linear)
- 5: Return strategy

It's a game agent vs env!

- Build the arena
- Play to win!

Theorem ([DeGiacomoRubinIJCAI18])

FOND for LTL_f/LDL_f goals is:

- EXPTIME-complete in the domain (assuming a logarithmic representation as in PDDL);
- 2EXPTIME-complete in the goal. *EXPTIME-complete for PPLTL*

Same as classic FOND

Note we have separated costs in the model (DOM) and the task GOAL!
(c.f. data vs query complexity in Databases)

Planning in Nondeterministic Domains for PPLTL Goals

FOND for LTL_f goals

1. Given a FOND domain D and an PPLTL formula φ
2. Compute the symbolic DFA A_φ for φ (linear)
 - using quoted subformulas as fluent
3. Encode the symbolic DFA A_φ in FOND (linear)
 - requires conditional effects
4. Plan for the propositional goal “ φ ” (exponential)
5. Return (strong/strong cyclic) plan.

It's a game agent vs env!

- Build the arena
- Play to win!

Theorem ([DeGiacomoRubinIJCAI18])

FOND for LTL_f / LDL_f goals is:


- EXPTIME-complete in the domain (assuming a logarithmic representation as in PDDL);
- 2EXPTIME-complete in the goal. EXPTIME-complete for PPLTL

Same as classic FOND

Note we have separated costs in the model (DOM) and the task GOAL!
(c.f. data vs query complexity in Databases)

DECLARATIVE ...
... TO PROCEDURAL ...
... TO GAMES

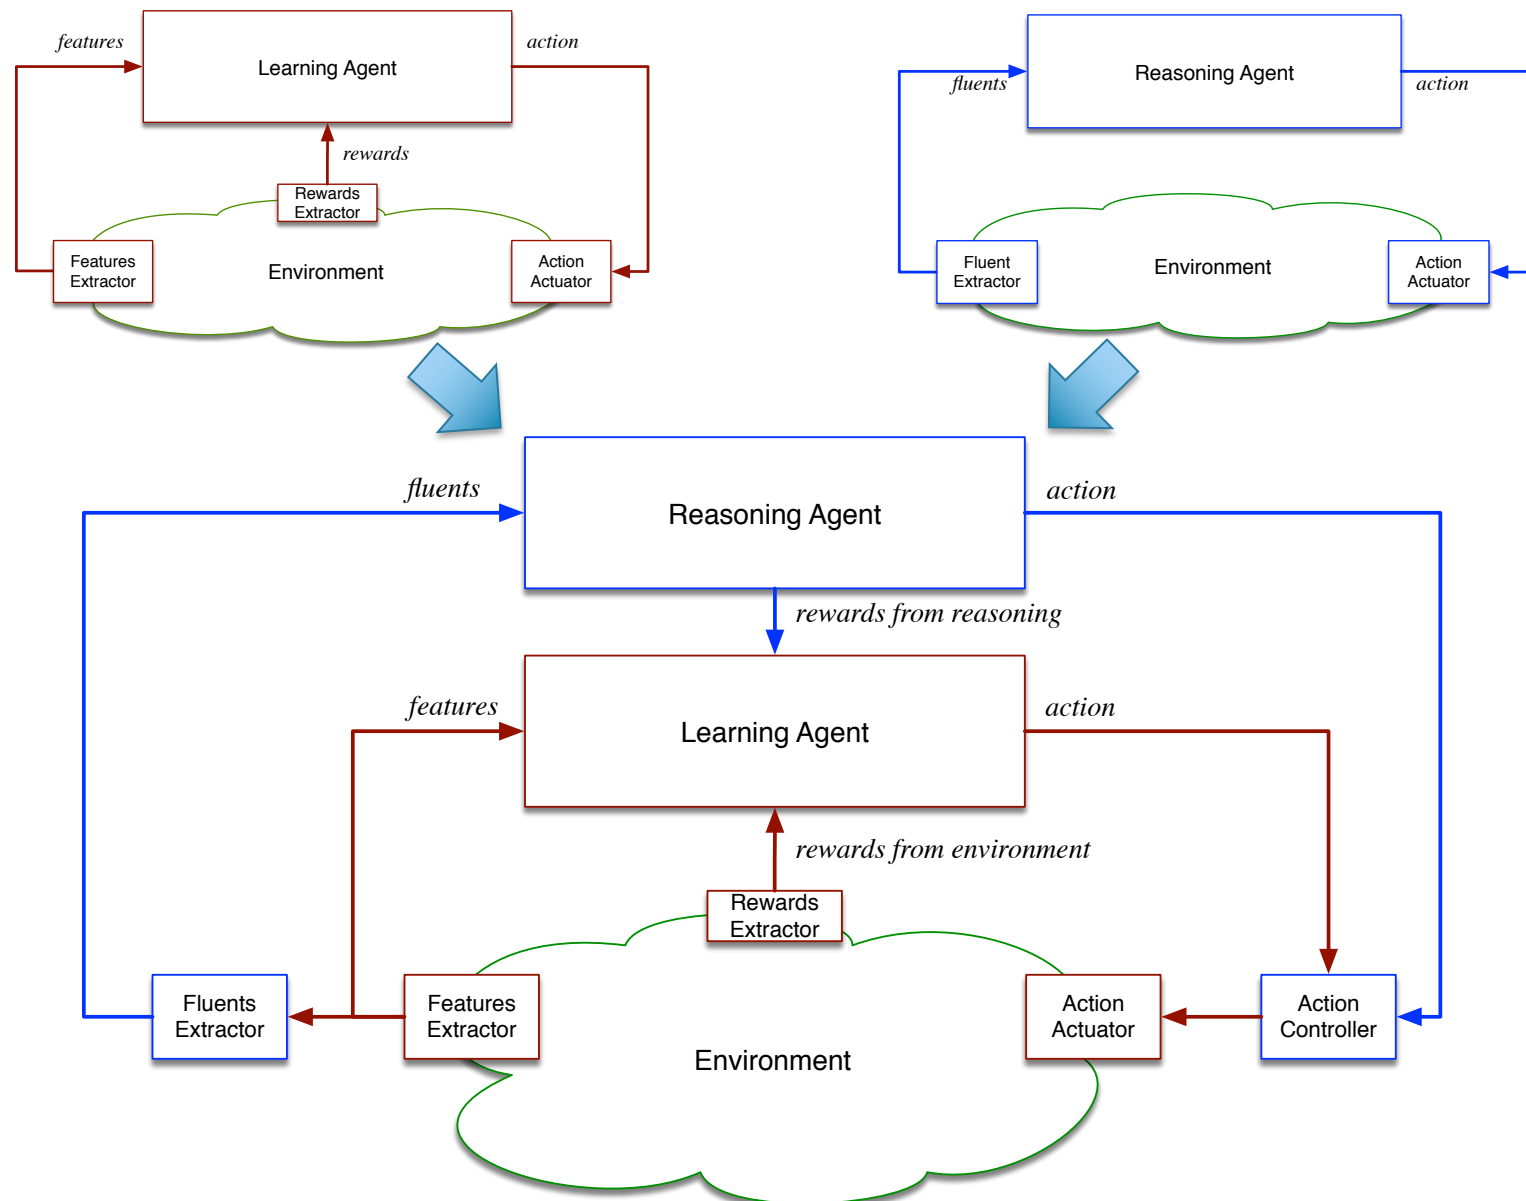
DECLARATIVE CONTROL IN RL



Combining Learning and Reasoning

Merging:

- Learning agent
 - Does reinforcement learning
 - Possibly deep reinforcement learning
- Reasoning agent
 - Does reasoning
 - Possibly on temporal specification as in formal methods



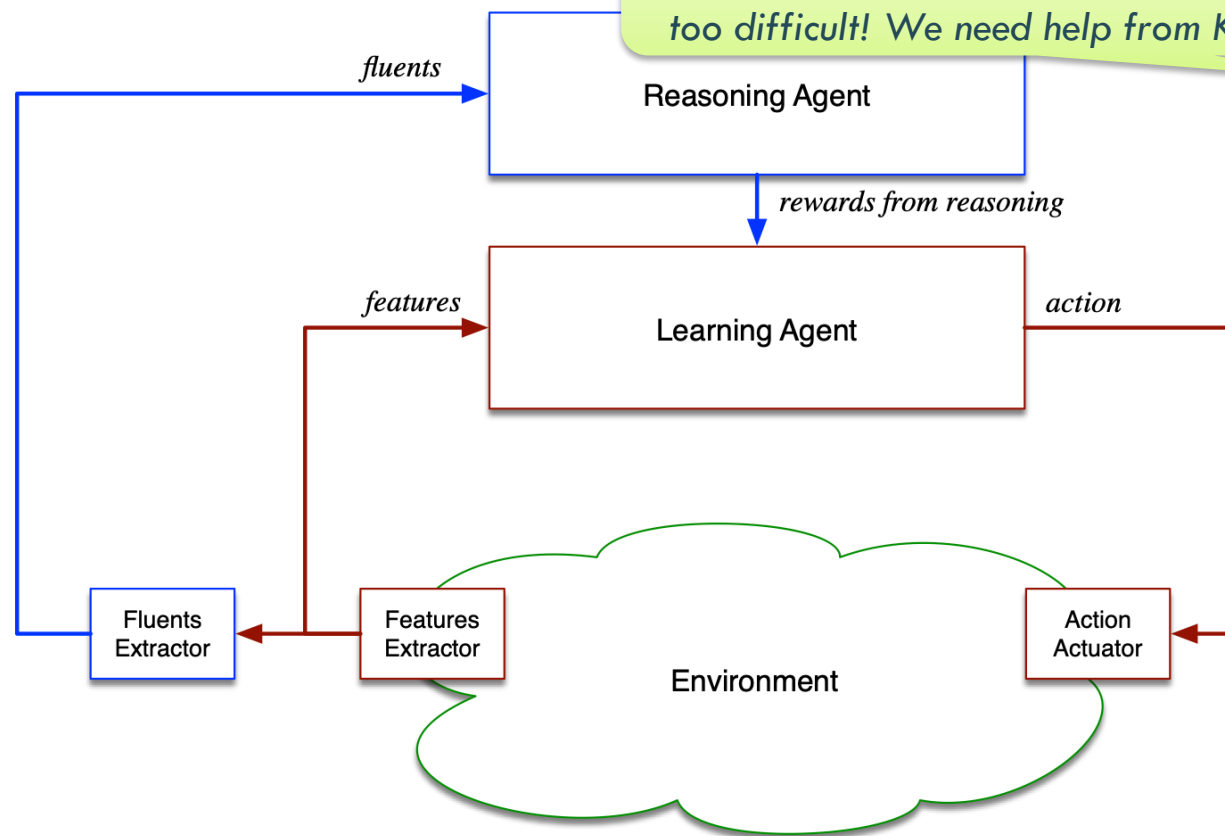
Reinforcement Learning with LTLf non-Markovian Rewards

Michael Littman @ IJCAI 2015:
"Coming up with rewards in MDPs is too difficult! We need help from KR!"

MDPs with non-Markovian rewards

- **Learning agent:** $\mathcal{M} = (S_{ag}, A_{ag}, Tr_{ag}, \cancel{R_{ag}})$
MDP without rewards
- **Reasoning agent:** $\mathcal{R} = (\mathcal{L}, \{(\varphi_i, r_i)\}_{i=1}^m)$
 φ_i in LTLf/LDLf $\longrightarrow \bar{R}_{ag} : (S_{ag}, A_{ag})^* \rightarrow \mathbb{R}$
non-Markovian rewards!
- **Mapping between S_{ag} and \mathcal{L}**

We can define equivalent MDP over an extended state space and do standard RL



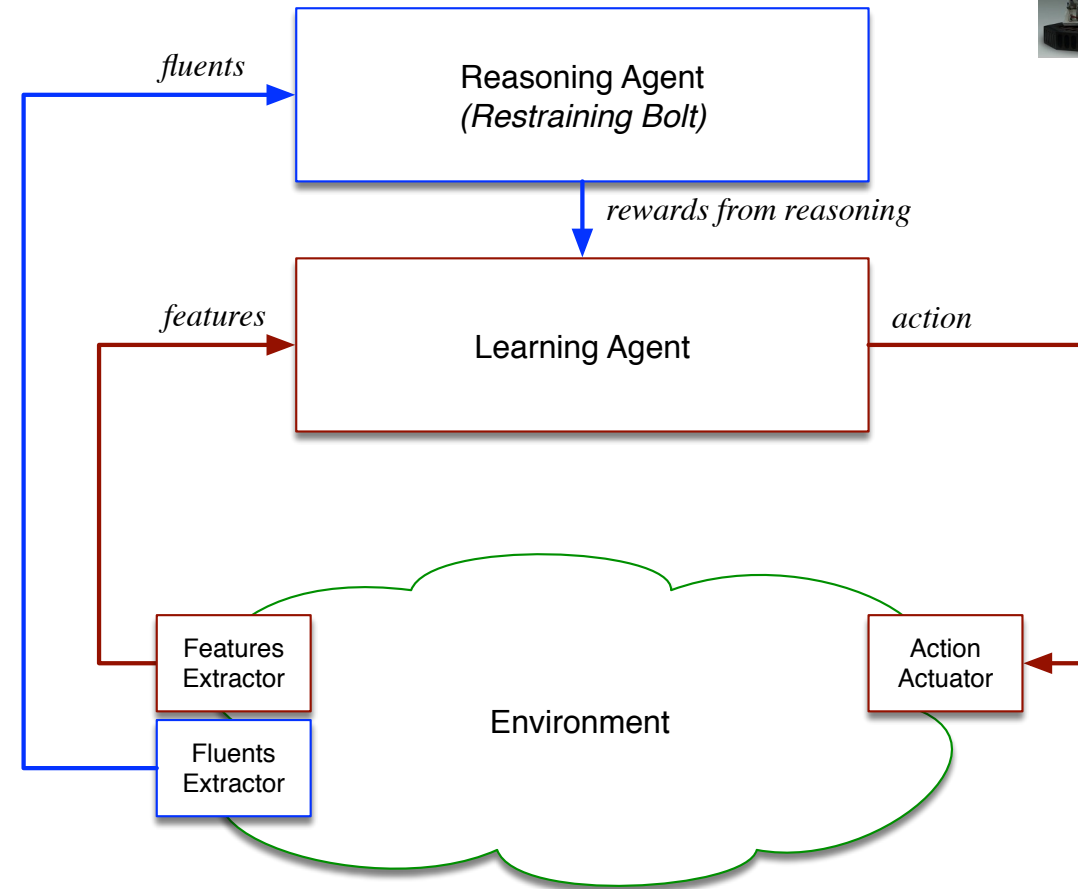
Restraining Bolts as Reasoning Agents




Double state representation (restraining bolts)

- Learning agent: $\mathcal{M} = (S_{ag}, A_{ag}, Tr_{ag}, \cancel{R_{ag}})$
MDP without rewards
- Reasoning agent: $\mathcal{R} = (\mathcal{L}, \{(\varphi_i, r_i)\}_{i=1}^m)$
 φ_i in LTLf/LDLf $\longrightarrow \bar{R}_{ag} : (S_{ag}, A_{ag})^* \rightarrow \mathbb{R}$
non-Markovian rewards!
- ~~Mapping between S_{ag} and \mathcal{L}~~

We can define equivalent MDP over an extended state space and do standard RL



Example: Breakout


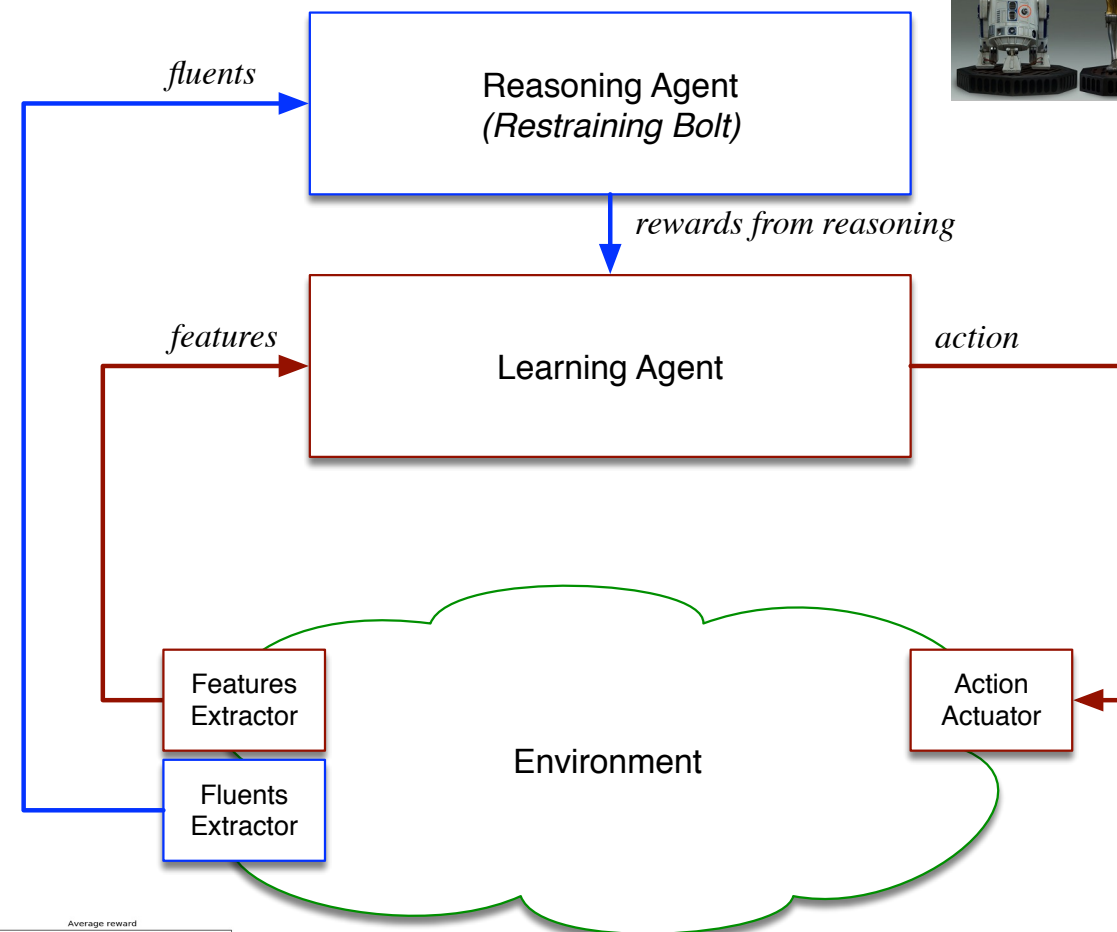


Learning Agent

- Features: paddle position, ball speed/position
- Actions: move the paddle
- ~~Rewards: reward when a brick is hit~~

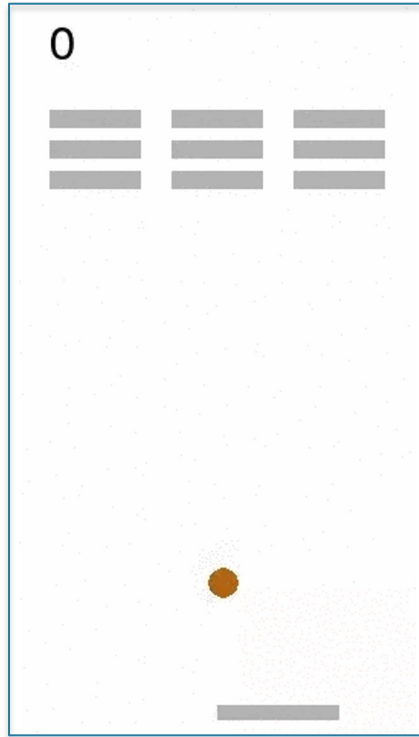
Restraining Bolt (Reasoning Agent)

- Rewards: break one column at the time left to right (all bricks in column i must be removed before completing any other column $j > i$)
- Fluents: bricks/columns status (broken/not broken)




<https://sites.google.com/diag.uniroma1.it/restraining-bolt>

The Agent Ignores the Fluents!



How the world is




How the agent sees
the world

Controlling RL with Logic for Safety

The idea of restraining bolt can be subscribed to that part of research generated by the urgency of providing **safety guarantees** to AI techniques based on learning.

- S. Russell, D. Dewey, and M. Tegmark. **Research priorities for robust and beneficial artificial intelligence**. AI Magazine, 36(4), 2015.
- ACM U.S. Public Policy Council and ACM Europe Policy Committee. **Statement on algorithmic transparency and accountability**. ACM, 2017.
- D. Hadfield-Menell, A. D. Dragan, P. Abbeel, and S. J. Russell. **The off-switch game**. In IJCAI 2017.
- D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mane. **Concrete problems in AI safety**. CoRR, abs/1606.06565, 2016.
- Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Konighofer, Scott Niekum, Ufuk Topcu: **Safe Reinforcement Learning via Shielding**. AAI 2018.
- Min Wen, Rüdiger Ehlers, Ufuk Topcu: **Correct-by-synthesis reinforcement learning with temporal logic constraints** IROS 2015.



Shields

Restraining Bolts as Shields for Frame

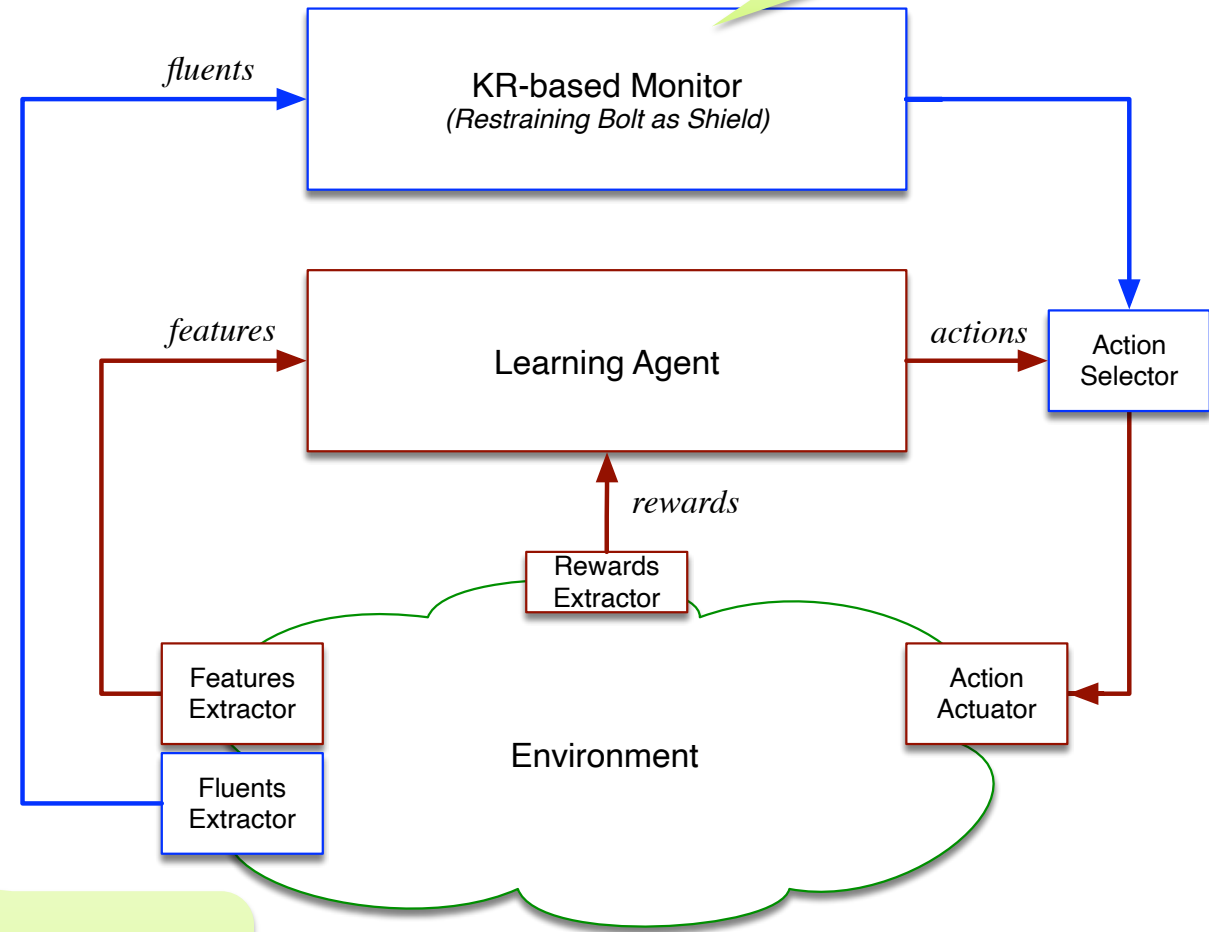
Poss(a,h) iff $h \models \varphi_{LTLf}$

Separate representations

- KR-based agent
- RL-based agent

Shielded execution

- KR-based agent acts as a monitor
- It disallows forbidden actions

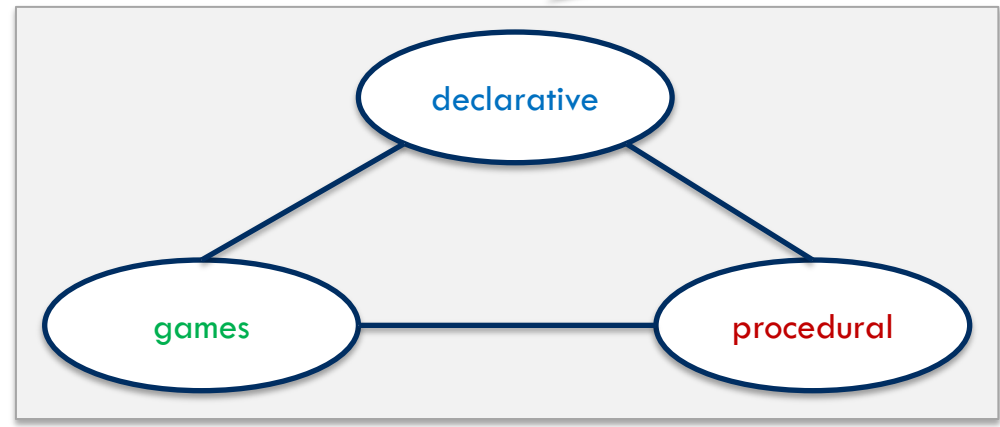


Very related to “*Framed Autonomy*” in ABPMS, see
AI-Augmented Business Process Management Systems: A Research Manifesto ACM Transaction on Management Information Systems, 2023

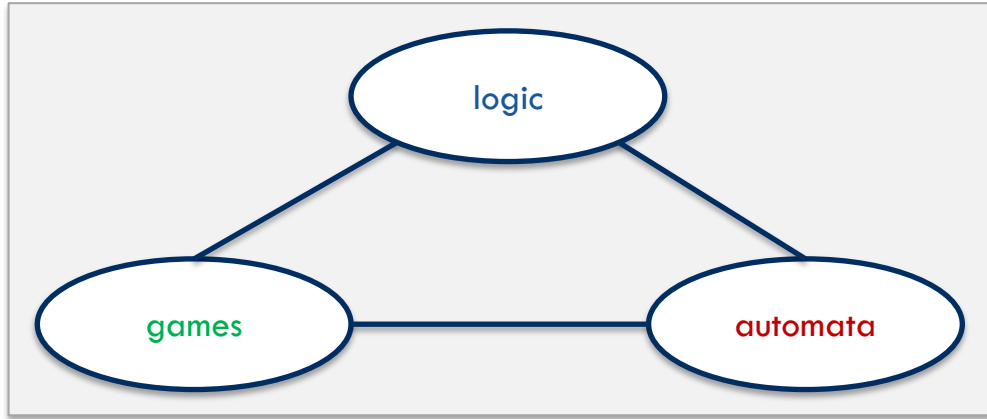
N. Alechina, M. Dastani, G. De Giacomo, B. Logan, G. Perelli, G. Varricchione. Pure-Past Action Masking, AAI2024

Conclusion

It has profound implication in how we can represent services



Powerful formal framework for controller synthesis



The instantiation on LTLf/PPLTL is particularly effective

