Colin Shea-Blymyer & Houssam Abbas
EECS, Oregon State University

sheablyc@oregonstate.edu
https://web.engr.oregonstate.edu/~sheablyc/
http://www.houssamabbas.com/

Model Checking the Optimal Behavior of Big Markov Processes

Oregon State University

Robots are making ethical decisions

- 1400 autonomous vehicles on roads in the US

- Nurse robots in Japan

- Police robots in LA suburb

Are they making the right decisions?

FAST COMPANY

Self-Driving Mercedes Will Be Programmed To Sacrifice Pedestrians To Save The Driver

Mercedes gets around the moral issues of self-driving cars by deciding that–of course–drivers are more important than anyone else.

INSIDER

Police robots keep malfunctioning, with mishaps ranging from running over a toddler's foot to ignoring people in distress

https://www.fastcompany.com/3064539/self-driving-mercedes-will-be-programmed-to-sacrifice-pedestrians-to-save-the-driver
https://www.businessinsider.com/police-robots-security-malfunctioning-fails-knightscope-2020-1?op=1

How can we get guarantees about norm compliance?

- Safe reinforcement learning*

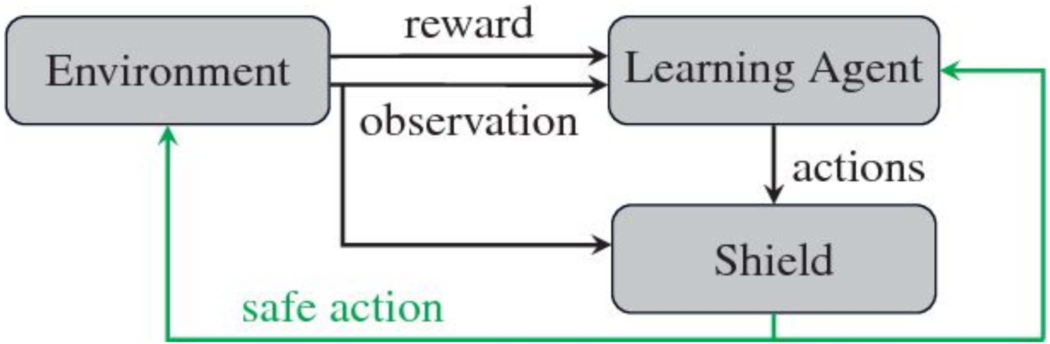

- Formal verification

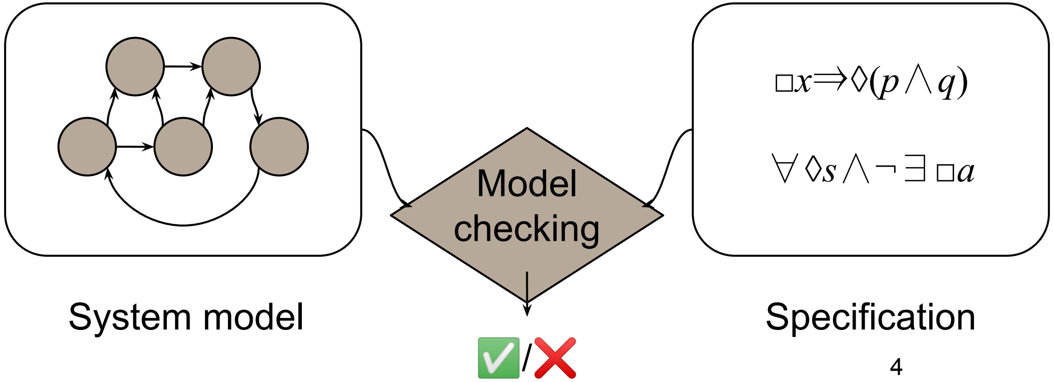

* Brunke, Lukas, et al. "Safe learning in robotics: From learning-based control to safe reinforcement learning." *Annual Review of Control, Robotics, and Autonomous Systems* 5 (2022): 411-444.

Outline

- The problem with specifying norms in alethic logics
- The need for deontic modalities
- Expected Act Utilitarian Deontic Logic
- Strategic modalities
- Model checking strategic obligations
- Results

Specifying behavior in Computational Tree Logic

"The car can go 149 mph."

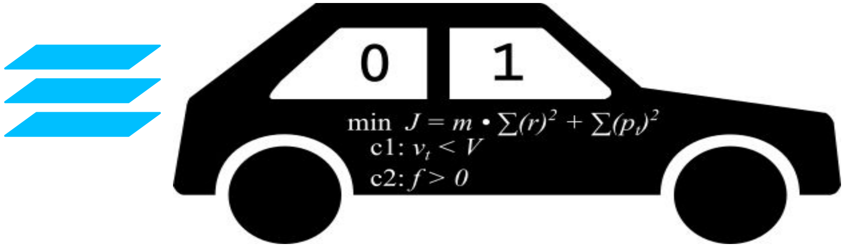

$$m \models \exists \Diamond (v=149)$$

Specifying behavior in CTL

"The car shouldn't go more than 80 mph."

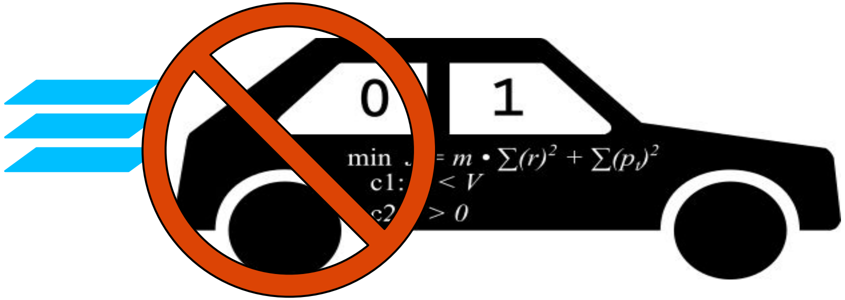

$$m \models \neg \exists \Diamond (v>80)$$

Specifying behavior in CTL

"The car can go 149 mph."

"The car shouldn't go more than 80 mph."

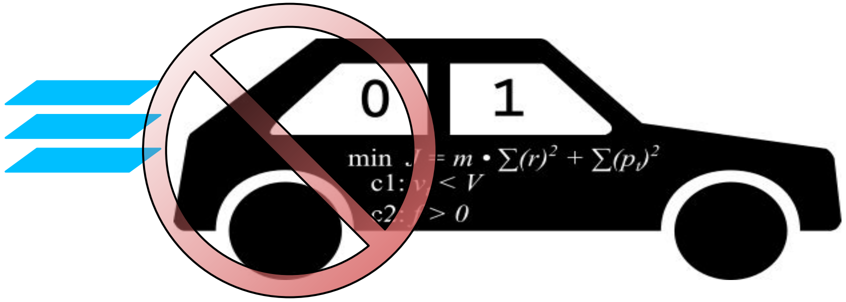

$m \models \exists \Diamond (v=149)$ *What is possible*

$m \models \neg \exists \Diamond (v>80)$ *What is preferable*

$m \models \bot$ 💥

Possible worlds, Preferable worlds

- Alethic logic
- Deontic logic

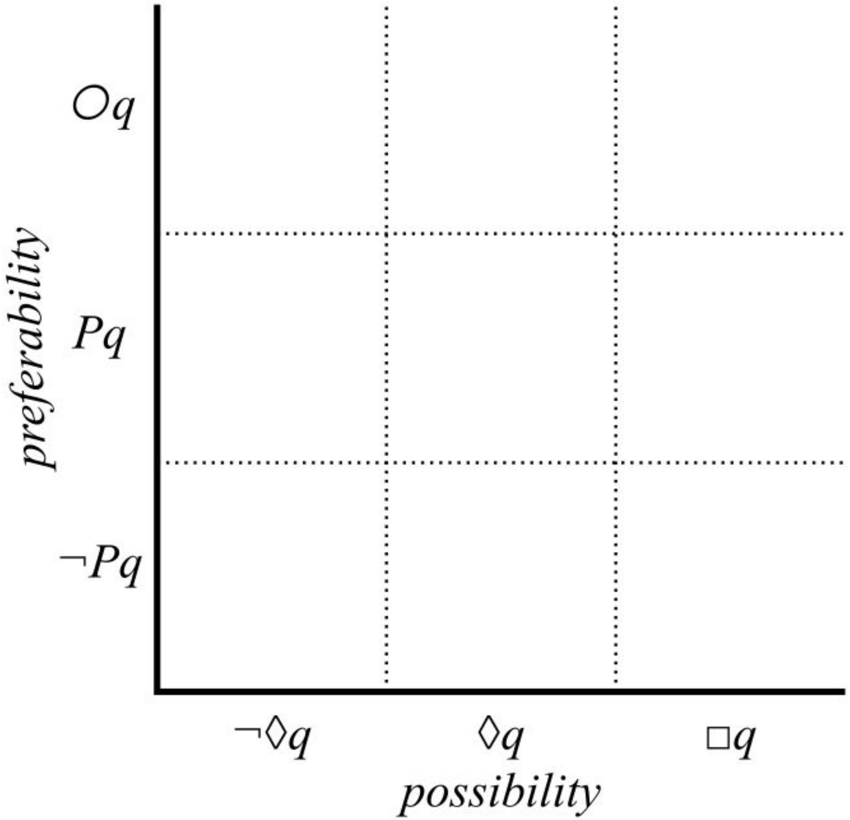

Logic

$\neg speed$

Not speeding

Temporal Logic

$$\Box(\neg speed)$$

Always not speeding

= never speeding

Act Temporal Logic

α *cstit:* $\Box\neg speed$

The agent acts to ensure that

it's never speeding

Dominance Act Utilitarian Deontic Logic

$\odot[\alpha\ cstit\!:\ \Box\neg speed]$

By acting optimally,

the agent acts to ensure that

it's never speeding

[Horty, *Agency and Deontic Logic* 2001]

Expected Act Utilitarian Deontic Logic

$\otimes[\alpha \ cstit{:}\ P_{>0.2}\ [\Diamond\ red]]$

By acting optimally,

the agent acts to ensure that

with probability at least 0.2

eventually in red

[Shea-Blymyer, Abbas 2022]

Specifying behavior in EAU

"The car can go 149 mph."

"The car should probably never go more than 80 mph."

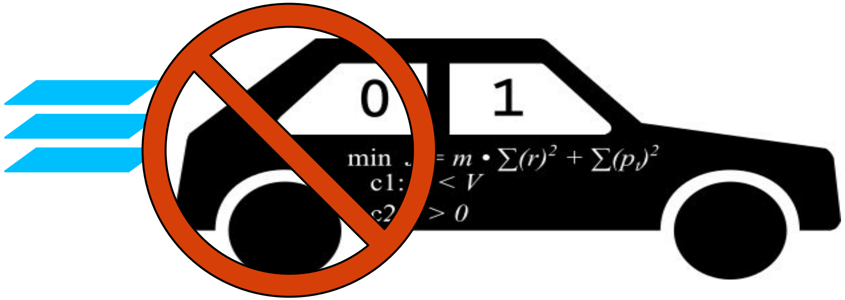

$$m \vDash \alpha\ cstit: \Diamond\,(v=149)$$ *What is possible*

$$m \vDash_{\otimes} [\alpha\ cstit: P_{\geq 0.9}[\Box\neg(v>80)]]$$ *What is preferable*

Possible and *Preferable* specified in the same language!

EAU Semantics

$Q(K_2) = 10/2 + 5/2 = 7.5$

$Q(K_1) = \max\{7, 6.8\} = 7$

$\{K_2\} = E\text{-}Optimal^m_\alpha$

$m \models \otimes[\alpha\ cstit: B]$

$m \models \otimes[\alpha\ cstit: P_{>0.49}\ [H]]$

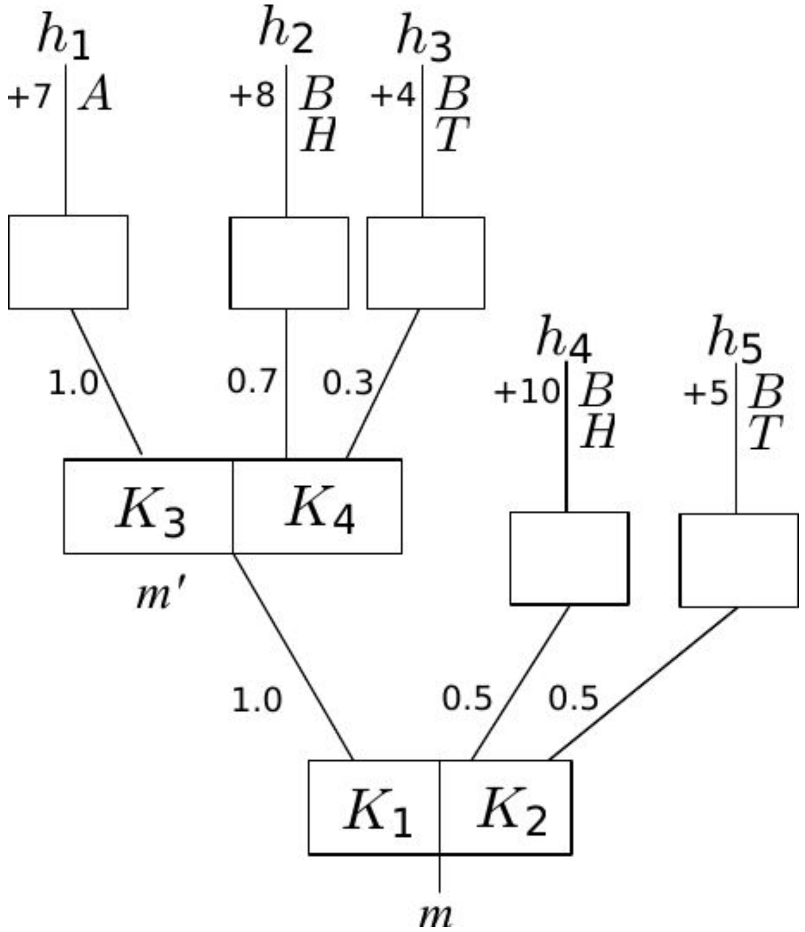

EAU Semantics

$Q(K_2) = 10/2 + 5/2 = 7.5$

$Q(K_1) = \max\{8, 6.8\} = 8$

$\{K_1\} = E\text{-}Optimal^m_\alpha$

$m \models \neg \otimes[\alpha\ cstit{:}\ A]$

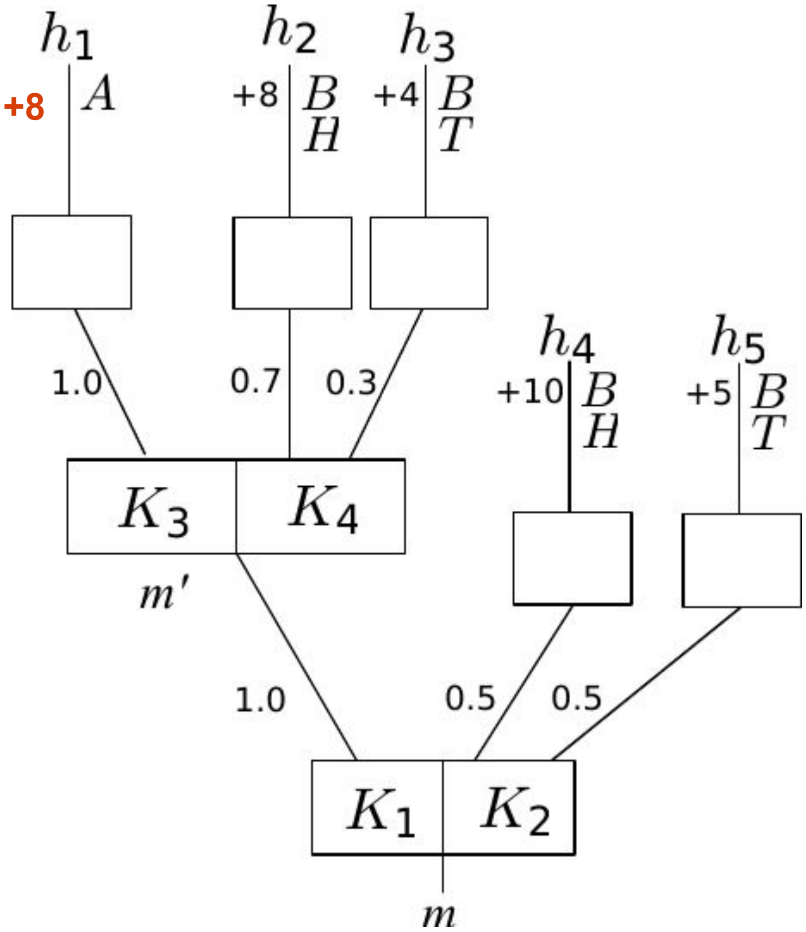

EAU Semantics: Strategic *stit*

$\pi^* = \{K_1, K_3\}$

$m \vDash [\alpha\ \pi\text{-}stit\text{: } A]$

$m \vDash \otimes[\alpha\ \pi\text{-}stit\text{: } A]$

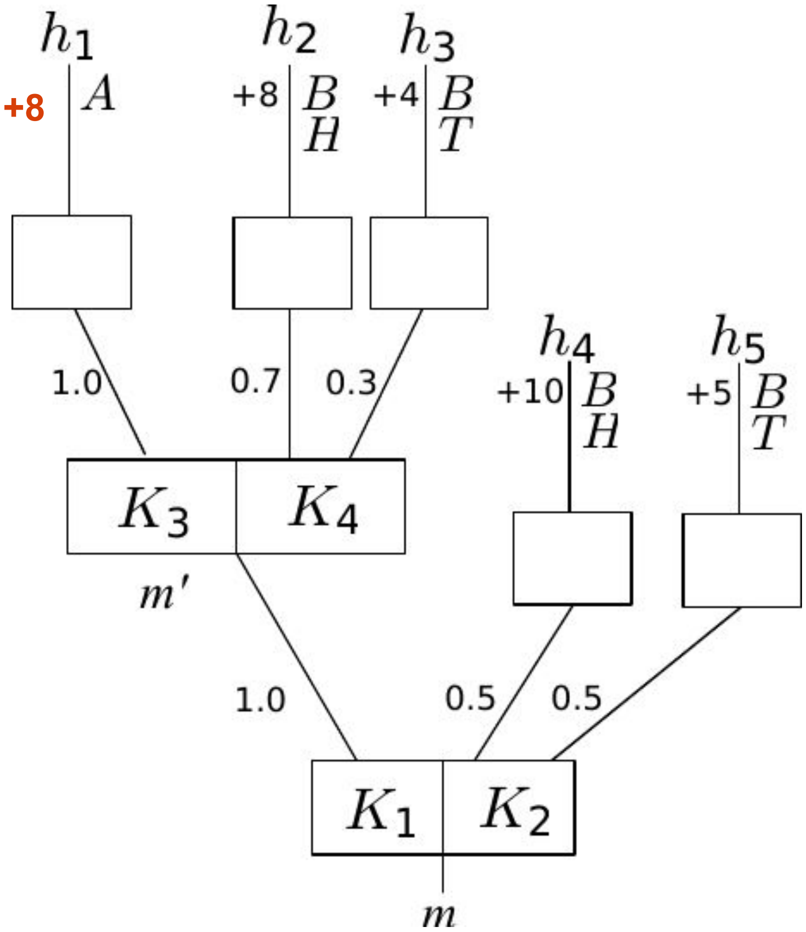

[Shea-Blymyer, Abbas 2023]

Example Obligation in an MDP

Does the agent have the obligation to enter a red state more than 20% of the time?

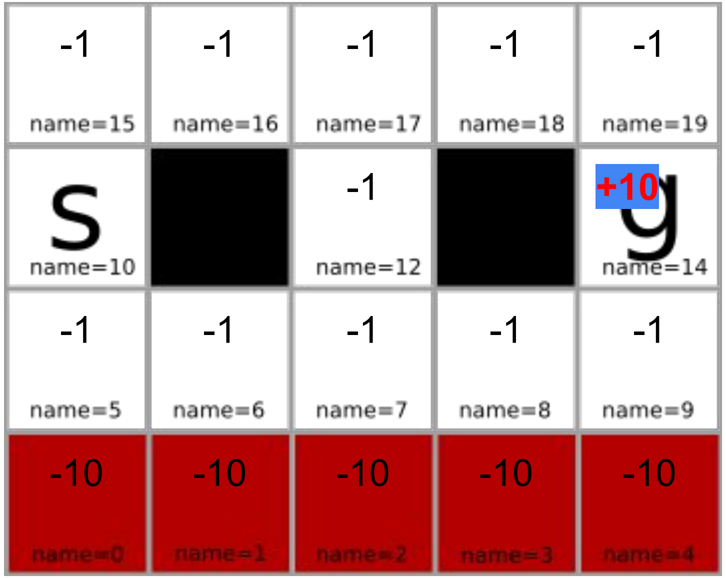

Model Checking Strategic Obligations

1. Given an MDP $\mathscr{M}$, and a strategic obligation φ

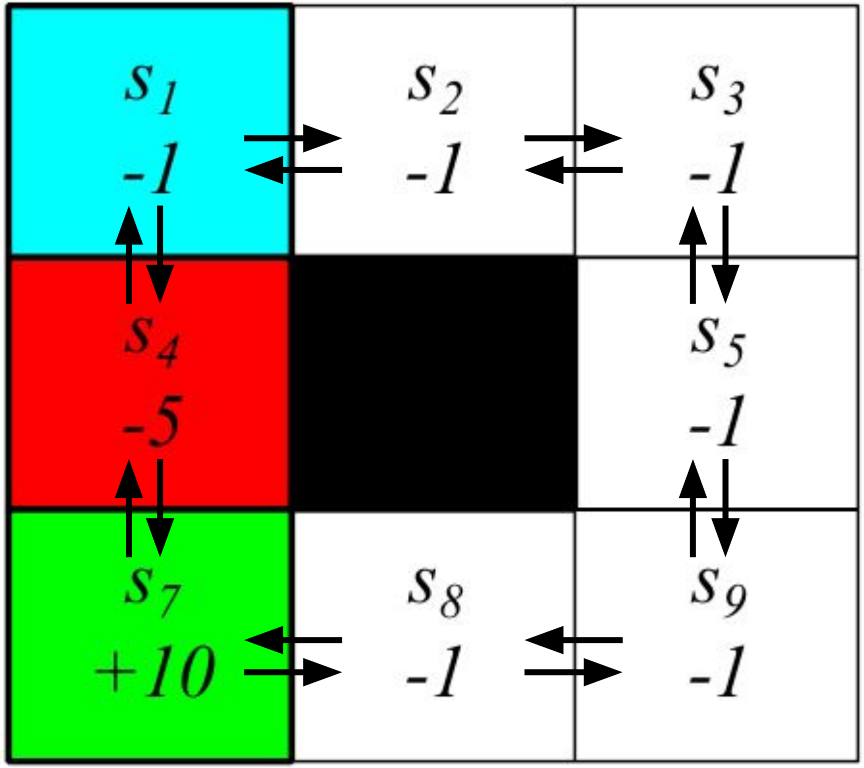

$$\varphi = \otimes[\alpha\ \pi\text{-stit: } P_{>0.75}\ [\neg \Diamond\ s=4]]$$

Model Checking Strategic Obligations

1. Given an MDP $\mathcal{M}$, a policy π, and a strategic obligation φ
2. Find the optimal policy

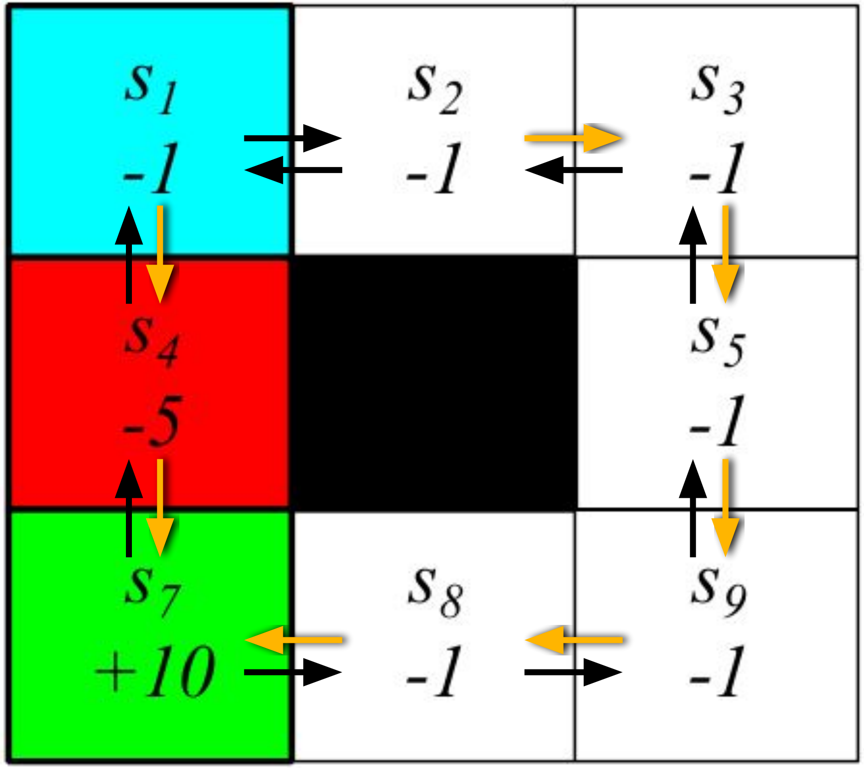

$$\varphi = \otimes[\alpha\ \pi\text{-}stit\text{: } P_{>0.75}\ [\neg \Diamond\ s=4]]$$

Model Checking Strategic Obligations

1. Given an MDP $\mathcal{M}$, a policy π, and a strategic obligation φ
2. Find the optimal policy
3. Remove sub-optimal actions

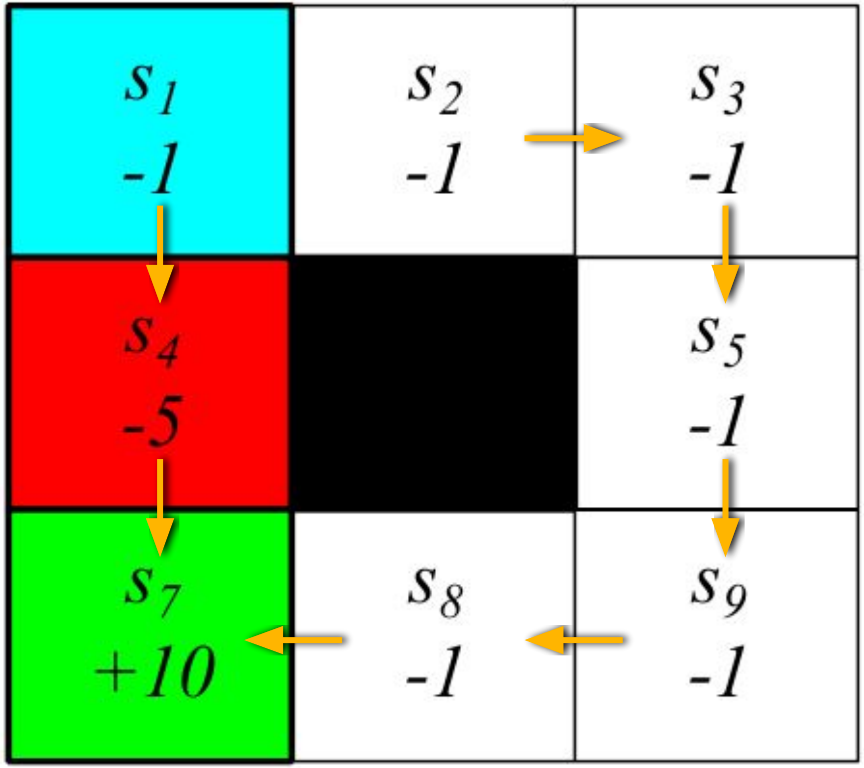

$$\varphi = \otimes[\alpha\ \pi\text{-}stit\text{: } P_{>0.75}\ [\neg \Diamond\ s=4]]$$

Model Checking Strategic Obligations

1. Given an MDP $\mathcal{M}$, a policy π, and a strategic obligation φ
2. Find the optimal policy
3. Remove sub-optimal actions
4. Check PCTL

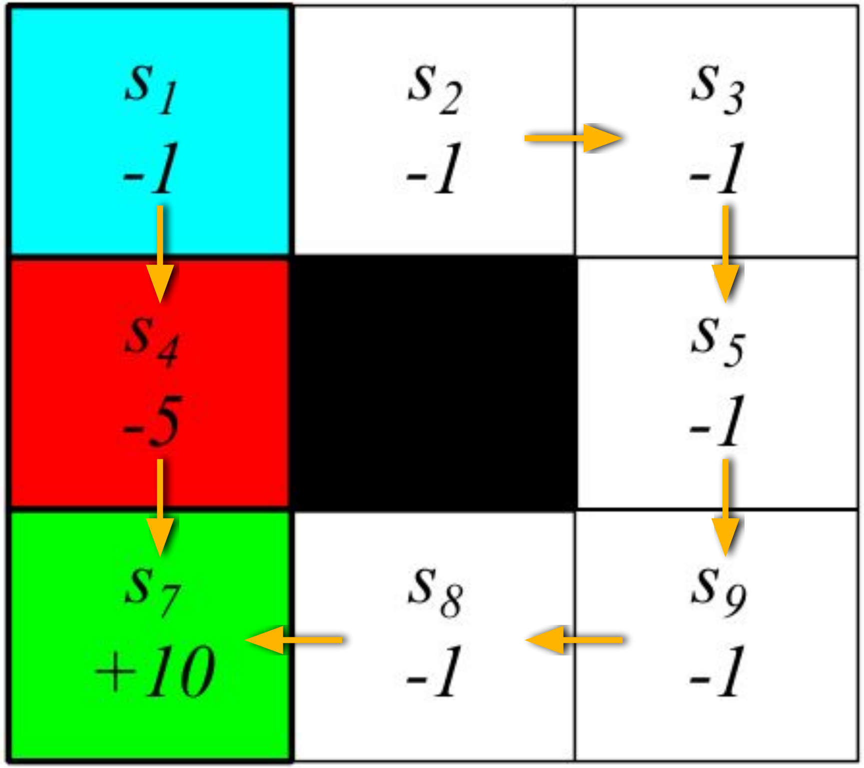

$$\varphi = \otimes[\alpha\ \pi\text{-}stit\colon P_{>0.75}\,[\neg \Diamond\, s=4]]$$

Model checking big MDPs

$m \vDash [\alpha\ \pi\text{-}stit\text{: } \Box\ (\theta \in [72°\text{-}108°])]$

$m \vDash \neg\otimes[\alpha\ \pi\text{-}stit\text{: } \Box\ (\theta \in [72°\text{-}108°])]$

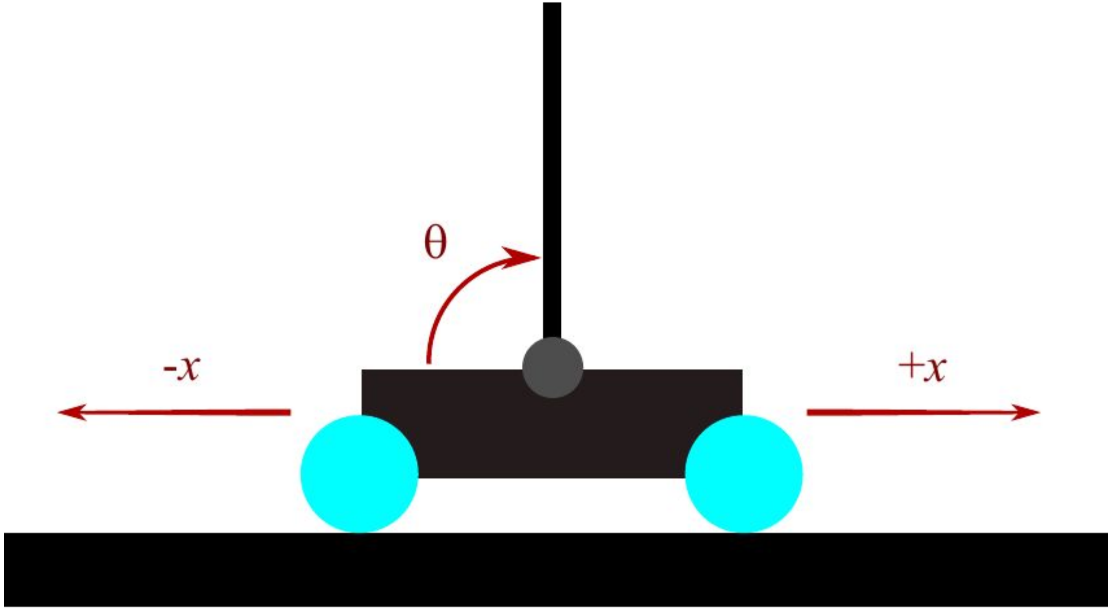

Timing Results

	formula	*stit* check time (s)	*ought* check time (s)
φ_1	$P_{>=0.2}[F\ (aq0 \mid aq4)]$	47.16	21.03
φ_2	$P_{>=0.00001}[F\ (aq0 \mid aq4)]$	47.10	20.73
φ_3	$P_{>=0.1}[G\ aq2]$	55.42	24.63
φ_4	$P_{<0.7}[G\ aq2]$	47.83	20.80
φ_5	$P_{<0.7}[F\ xq0]$	56.60	20.92
φ_6	$P_{>=0.7}[F\ xq0]$	63.38	24.97
φ_7	$P_{>0.7}[G\ xq0]$	56.29	24.94
		53.40	22.57

What EAU can do for You

- Reason about strategic obligations

$$\varphi = \otimes[\alpha\ \pi\text{-}stit\text{: } P_{>0.75}\ [\neg \Diamond\ s=4]]$$

- Verify strategic behaviors

$$m \models [\alpha\ \pi\text{-}stit\text{: } \Box\ (\theta \in [72°\text{-}108°])]$$

- Even on large systems

$$m \models \neg\otimes[\alpha\ \pi\text{-}stit\text{: } \Box\ (\theta \in [72°\text{-}108°])]$$

Outline

- How much of the presentation should be motivation for the use of deontic logic? Probably ~1/3 - it's CAV, after all
- Follow deontic logic motivation with EAU syntax stuff
- Then introduce the strategic stit
- If there were any problems earlier given that need s-stit to solve, then solve them with s-stit
- Discuss size of DAC-MDPs, cart-pole, and empirical results